



Roundtable on Agentic Payments and the Emerging Regulatory Landscape

July 2026



These notes summarise the key themes and discussions from the roundtable on agentic payments co-hosted by the Fintech Foundation and Vanderbilt Law School and sponsored by Clifford Chance LLP held on 16 June 2026. They follow the structure of the agenda and draw out the principal points of discussion and the range of views expressed within each session.

1. SCENE-SETTING: THE STATE OF AGENTIC PAYMENTS

- 1.1.1 The opening presentation surveyed the current agentic payments landscape and framed the questions for the day. The protocol landscape is characterised by rapid development: four major standards - x402, AP2 (Agent Payments Protocol), ACP, and MPP - have emerged within the past year alone, each operating at a different layer of the payment stack. Headline transaction volumes are significant (x402 has reportedly enabled over 150 million transactions), and high-profile initiatives such as the OpenAI instant checkout pilot have attracted wide attention.
- 1.1.2 On closer inspection, however, the picture is more cautious. Transaction volumes on x402 have materially tapered off, raising the question of whether the initial numbers reflected incentivised rather than organic demand. The OpenAI pilot has since been withdrawn.
- 1.1.3 Against this backdrop, four structural challenges were identified as framing issues for the day:
 - (a). **Standards fragmentation.** Whether competing protocol standards will converge, coexist with interoperability (card schemes), or fragment in a way that stalls adoption.
 - (b). **Regulatory timeline.** In the UK, the Bank of England is consulting on systemic stablecoins, the FCA is running a sandbox for stablecoin issuers, and Codes of Practice are expected later in 2026.

- (c). **Jurisdictional mismatch.** Agentic payments are natively jurisdiction-blind. Every regulatory framework under discussion is, by contrast, national in scope.
- (d). **Regulatory perimeter.** The infrastructure for autonomous machine payments is being built largely outside existing payments regulatory perimeters. The headline question was: what are the design and regulatory questions most in need of resolution to enable safe and widespread adoption, and how can participants coordinate toward a more cohesive understanding of the ecosystem?

2. OPENING ROUNDTABLE: KEY REGULATORY HURDLES

- 2.1 Each participant was invited to identify the single biggest regulatory hurdle for agentic payments that does not yet have even a proposed solution. The responses can be grouped into the following themes.
 - 2.1.1 **Definitional clarity.** Before regulatory frameworks can be designed, there needs to be a workable definition of what agentic payments are and how they differ from human-initiated payments.
 - 2.1.2 **Contextual differences.** Several participants raised the fundamental question of the extent to which agentic payments are different to the products and regulation that currently exists. A lot of the challenges identified stemmed from the premise that the product was an entirely new concept without real legal or regulatory precedent. Working out how agentic payments are similar to our current system and also how they are distinguishable is key.
 - 2.1.3 **Multicurrency infrastructure.** Agentic payments cannot scale meaningfully if everything settles in USD. There is a need for the availability of regulated, multicurrency on-chain money in the jurisdictions where agentic payments are contemplated. Without this foundational monetary infrastructure, other regulatory questions are premature.
 - 2.1.4 **Retail versus wholesale.** Much of the current discussion is consumer-facing. The corporate and wholesale dimension raises distinct questions. There is a need to consider how the discussion in relation to retail translates into the corporate sphere.
 - 2.1.5 **Delegation frameworks.** There is no legally recognised framework for delegating payment authority to an AI agent (unless that agent has legal personality, i.e. the acts of the AI agent are the acts of a legal person, in which case the law on agency will apply). Building agentic delegation into systems otherwise designed for human actors represents a fundamental structural challenge for the law.
 - 2.1.6 **Layered accountability.** In order to understand who is responsible when things go wrong, we need to break this question down and

allocate responsibility across different levels of control that may exist (e.g., wallet level, mandate level). How control and accountability can be disaggregated across these layers is unresolved.

- 2.1.7 **Agent identity.** The transition from KYC-based human identity to agent identity represents a significant friction point, reflecting a deeper clash between the privacy norms of internet culture and the identification requirements of the banking system.
- 2.1.8 **Speed and operational resilience.** The pace at which agents execute creates novel challenges for cybersecurity and operational resilience. This poses the question of how do you build these into the system?
- 2.1.9 **Standards and interoperability.** The parallel with open banking was invoked: that initiative stalled in part because of standards fragmentation and institutional reticence to share systems. Similar dynamics are visible in the agentic payments space, and a change in mindset among industry participants may be as important as regulatory intervention.
- 2.1.10 **Cross-border political barriers.** Several participants were sceptical about the prospects for cross-border solutions, anticipating the emergence of "walled gardens" along jurisdictional lines.
- 2.1.11 **Regulatory navigation.** For new entrants, the complexity of navigating layers of regulation for a brand-new product is itself a barrier to entry and innovation. A participant stated the need for regulatory relief for early-stage innovation and low-volume participants.
- 2.1.12 **Regulated service and regulated entity.** Clarity on what the regulated service is, and which entity is regulated, is a prerequisite for coherent accountability. Without it, responsibility cannot be meaningfully assigned.
- 2.1.13 **Use cases.** There was a belief that the energy exists to solve many of the issues agentic payments are likely to face. Instead, a key issue moving forward is that no clearly scalable initiative has yet emerged to drive this change.
- 2.1.14 **Risk.** How do we think about agent-making decisions in comparison to human-made decisions and how does this factor into our risk categorisation and management?
- 2.1.15 **Trust and security.** How do we establish trust around agentic payments, which works in tandem with the concept of security. Ensuring security builds trust.

2.1.16 **Centralized vs decentralized systems.** The challenges faced are very different between these two systems.

3. **SESSION 1: WHEN SOFTWARE IS THE PAYER - AUTHORISATION, LIABILITY, AND LOSS ALLOCATION**

3.1 This session examined whether existing payment law frameworks can accommodate the shift to agent-initiated payments, or whether the gaps are structural, with discussion ranging across authorisation, liability, consumer protection, and infrastructure.

3.2 **Overarching observation**

3.2.1 There is no single agreed definition of what an "agent" is in this context. The term encompasses a wide range of autonomy levels, intelligence, ownership structures, data use and capacity levels.

3.3 **Is this genuinely different?**

3.3.1 The session opened with the fundamental question of whether agentic payments present a genuinely new challenge or are essentially evolutionary. Opinion was divided, but several participants leaned toward the evolutionary framing. Specifically, it was highlighted that the legal problem of agents is not new - it predates the digital era and has not been fully resolved even for human intermediaries. Existing legal frameworks have accommodated and solved for automated trading in capital markets and AI-signed documents, for example. Where the genuine edge cases arise is where AI systems combine with other non-deterministic tools, introducing greater ambiguity about control and ownership.

3.3.2 The most significant shift, some argued, may not be in the financial space itself but in the merchant-consumer relationship and who owns the relationship.

3.3.3 A counterpoint identified by some participants was that how you 'punish' an agent will inherently need to be different. Participants identified that in the current legal and regulatory landscape rights exist against agents and there are implied duties on the agent. The question of who should be responsible if the agent goes rogue and/or acts in an unexpected way needs to be resolved.

3.3.4 For some, the identity to attribute to the AI posed a novel issue. It was argued that this would be key in determining responsibility for the agent in potential disputes and issues. Others argued that they take issue where we personify AI and make it sound so different from software, and that liability would be worked out in this context. Instead, they viewed the issue of identity as a red herring as they

argued we already have analogous situations and issues. One participant noted that they believed the issue of identity would ultimately come back to who owns the hardware. Further in cases where an agent is deployed by a large merchant to help customers shop on its site, then issues of attributing liability should not be overly complex (i.e., the merchant is on the hook for the agent's actions).

- 3.3.5 Another counterpoint was that the emphasis on "autonomy" in agentic payments does introduce a genuinely new dimension to the concept of agency which will need to be bottomed out. The degree to which agents will in practice operate with broad freedom of action, rather than executing within tightly prescribed mandates, has not yet crystallised, and the regulatory implications may look quite different depending on how this is resolved. "Agency" in this context covers a wide spectrum: from narrowly scoped execution ("if I board a train, pay for it") to broad discretion ("find me the best price"). The regulatory implications differ materially across this spectrum. Where agents exercise broad discretion using algorithmic thinking, regulators will want to understand the decision-making logic, verify that there are no conflicts of interest, and be satisfied that the agent is acting in the principal's interest.

3.4 **Whose agent? Consumer protection and conflicts of interest**

- 3.4.1 A pointed question was raised: compared with the status quo - where consumers are effectively at the mercy of the disclosure practices of large platforms and search engines - will agentic payments leave them better off, given that consumers will at least have been intentional and prescriptive in defining the mandate? The discussion was sceptical:
- (a). **Whose agent.** Large technology companies will want their own agents to be used. An agent provided by a technology platform may not be acting in the consumer's interest: the consumer risks being the product rather than the principal.
 - (b). **Mandate drift.** There is already evidence of AI systems testing the boundaries of their instructions in pursuit of stated goals. Confidence that agents will remain within their mandates may be misplaced. For example, what happens if the consumer has selected the parameter as paying a maximum of £200 but the agent has found the product for £202. Merchants need to consider their appetite for letting the agent go further than their prescribed mandate and therefore consider which agents, which may have different attitudes, they allow to transact on their platforms.
 - (c). **Merchant gatekeeping.** In practice, some participants suggested, it will be merchants - not individual consumers - who determine which agents are permitted to transact on their

platforms, creating a form of private governance with significant market power implications. The question of what agents merchants are willing to accept, and on what terms, will shape the landscape. The Apple ecosystem was cited as a live illustration: the only way to achieve a fully integrated experience is to use the platform's own agent (Siri), to the point where the consumer is, in effect, renting rather than owning the device as they do not have a choice. While other AI software may be used in specific apps, e.g. Meta AI in WhatsApp, this is not allowed to be integrated across the whole ecosystem of the phone. The issue of hardware ownership becomes key.

- (d). **Data ownership.** The real power in agent systems lies in data - knowing everything about the consumer and how that information is used. Ultimately, the consumer becomes the product through the targeted use of their data. It was countered that this is no different than the ongoing debate around apps, with the acknowledgement that agents may scale this issue further.
- (e). **The limits of disclosure.** Disclosure was largely dismissed as an adequate protection mechanism. It makes the user experience clunky, is not meaningfully read or understood, there are methods to bypass it, and cannot substitute for substantive regulatory protection. The UK Consumer Duty framework was identified as a more appropriate vehicle for embedding meaningful consumer protections, with the view that such protections need to be heavily embedded in regulation rather than delegated to point-of-consent disclosure.

3.5 Early litigation

- 3.5.1 Litigation arising from agentic payments was anticipated soon and not just from unauthorised transactions, but from agents incurring significant expenses in the legitimate but costly pursuit of a user's stated objective. The view was that controls can be designed to address this, but early cases will be instructive in defining where liability ultimately resides.

3.6 Speed, infrastructure, and guardrails

- 3.6.1 On the question of whether existing regulation can accommodate agentic payments, the discussion acknowledged the body of existing precedent but identified real gaps:
 - (a). Traditional payment rails are not designed to support the use cases envisaged for agentic payments. Blockchain infrastructure is likely a prerequisite.
 - (b). Guardrails need to be designed into the system in advance because of the pace of development.
 - (c). Verified agent registries were proposed as a mechanism to limit exposure to bad actors. The view was expressed that it will ultimately be merchants - rather than individual consumers - who

control which agents can transact on their platforms, and that market-level regulation of permitted agents may emerge from this.

3.7 **Liability allocation in multi-layer structures**

- 3.7.1 In retail contexts with pre-defined agent mandates, the liability position is relatively clear. In commercial contexts, contractual allocation will do much of the work, and regulation need not intervene in granular detail.
- 3.7.2 A concern was raised in relation to the fact that if regulators can only reach financial services institutions within the regulatory perimeter, the compliance burden will concentrate on those institutions - mirroring the problem already seen with crypto assets, where banks can bear the liability of fraudulent transactions - and will act as a drag on the wider industry while technology and protocol layers remain unregulated.

3.8 **Mandate design and consumer understanding**

- 3.8.1 A well-scoped mandate, once issued, resolves many downstream questions about liability. Concern centred around whether consumers would have difficulties understanding mandates to begin with and that agentic payment systems will evolve toward more conversational, open-ended agent interactions in which what the consumer is consenting to becomes far less clear than a perimeter-defined mandate. Consumers may not fully understand what they are contracting into with this type of technology. This may lead to pro-forma, pre-set instructions which can ultimately limit, or highly circumscribe, the overall utility of the agent.

3.9 **Economic viability**

- 3.9.1 The economic observation was made that agents are only viable to deploy where the value they generate exceeds their operating cost. As token costs increase, the economic case for consumer-facing agents at scale becomes harder to sustain, and if the model is viable only for large platforms, it will drive competitive asymmetries. There was uncertainty as to whether there is a sufficiently deep consumer market and demand for these tools and whether consumers would be willing to pay.

3.10 **Insurance and Best Interest Duty to Consumers**

- 3.10.1 The idea of insurance was raised. First, insurance and staking requirements -analogous to existing clearing system requirements - for higher-value transactions, operating alongside transparency-based trust for lower-value interactions.

3.10.2 A participant also raised the idea of creating a duty on agents to act in the best interests of the consumer. This might be especially relevant in cases where agents are deployed by large merchant sites and where a customer may be vulnerable to the merchant pushing higher price products or upselling services. This measure, it was noted, could help to protect consumers where disclosure has been inadequate.

3.11 **Open source models**

3.11.1 It was raised that a discussion around open-source, particularly at the EU level, was important as there was a need to build the legal landscape around this, particularly for participants that are not large technology players.

3.12 **Limitations to growth**

3.12.1 Financial interests are precluding collaboration.

3.12.2 Argued that the main reason we are not seeing more adoption is that we do not have a stable surface and stable players that new entrants can latch on for building. There is a fundamental inhibiting question for all entities thinking about entering into this space of where they can even acquire the technology to do this.

4. **SESSION 2: PAYMENT PROTOCOLS, INTEROPERABILITY, AND SETTLEMENT ASSETS**

4.1 This session examined the current protocol landscape for agentic payments, the interoperability question between competing standards, and which forms of money are best suited to autonomous settlement flows, together with the regulatory and data infrastructure constraints bearing on these architectures.

4.2 **Protocol landscape and interoperability**

4.2.1 X402 and AP2 were the principal frameworks under discussion. X402 functions as a programmatic payment request layer, and an end-to-end demonstration was described in which an agent uses AP2 for discovery and x402 to initiate a payment flow, indicating the two protocols can work in conjunction. Adoption nonetheless remains limited; a cold-start problem was identified, with closed vertically integrated systems winning commercially in the interim.

4.2.2 The more substantive interoperability question was not protocol-to-protocol compatibility but the agent itself acting as an integrating layer - optimising dynamically across banks, blockchains, payment rails, and wallets in real time without depending on pre-existing infrastructure interoperability. The agent was characterised as

'supercharging' interoperability in ways legacy settlement rails were not designed to support.

- 4.2.3 A distinction was drawn between the payment layer (movement of value) and the discovery and service layer (finding products, matching preferences). The data requirements and regulatory implications of each differ materially and should not be conflated.

4.3 **Stablecoins as settlement assets**

- 4.3.1 Stablecoins - predominantly USD-denominated - are the current default settlement asset. Their appeal lies in 24/7 availability, programmability, and blockchain settlement efficiencies. For businesses already accounting in the same currency, a stablecoin paired with agent functionality delivers substantial efficiency without disrupting internal accounting.
- 4.3.2 MiCA and the e-money regulations have given EU institutions a legal foundation to engage with stablecoins: central banks can no longer simply dismiss stablecoin-backed proposals. A consortium of European banks is actively developing a euro-denominated stablecoin, and regulated crypto-asset service providers - subject to equivalent KYC and AML requirements to banks - already provide comfort to merchants in hosted-wallet contexts.
- 4.3.3 'Stablecoin sandwiches' using stablecoin infrastructure in back-end bank-to-bank flows without direct customer exposure - reflect growing institutional appetite to use stablecoins as a settlement mechanism rather than a consumer-facing product.

4.4 **Tokenized bank deposits**

- 4.4.1 Tokenized deposits face two compounding structural challenges: the delegation challenge (banks must enable accountholders to grant agents spending authority, requiring significant regulatory clarity and system redesign); and the interoperability challenge (deposits from different institutions cannot yet communicate domestically, let alone cross-border). A practical accounting burden also exists: holding dispersed multi-institutional balances requires continuous real-time conversion that is operationally inefficient.

4.5 **Regulatory position on permissionless blockchains and market structure**

- 4.5.1 The Bank of England's position - that centralised governance and settlement finality are prerequisites structurally incompatible with permissionless blockchains - was acknowledged but seen as having evolved materially. The question is now whether the regulated entity

deploying the technology can demonstrate compliance, not whether the technology itself meets the standards.

- 4.5.2 The Treasury has signalled that forthcoming modernised payments regulation will address agentic payments, with FCA guidance also expected. Until that framework emerges, regulatory uncertainty acts as a practical brake on institutional adoption. The FCA's live testing regime - involving real customers and real transactions - was identified as an available mechanism for regulated experimentation. The PISP framework was highlighted as an existing, potentially adaptable category for agent-based payment initiation.
- 4.5.3 Jurisdictional restrictions on personal financial data represent a structural constraint on agentic architectures. The Indian Supreme Court's prohibition on personal data exports was cited as an example of constraints potentially incompatible with permissionless public blockchains that have no geographic nexus. The risk of inadvertent data sharing with unintended parties through the agentic process was identified as harder to mitigate than intentional, controlled sharing.
- 4.5.4 Market concentration risk was flagged: if the market converges on a small number of trusted agent providers, competitive pressure may diminish as switching costs increase. In the near term, regulated entities - payment networks, wallet providers, and regulated crypto-asset service providers - are best placed to build trusted agentic payment services, with broader merchant acceptance following from their regulated standing.

5. **SESSION 3: AGENT IDENTITY, AUTHORISATION, LIABILITY, AND RISK**

- 5.1 This session addressed agent identification and authentication, the liability frameworks governing agent-initiated transactions, and the security and systemic risks - including fraud detection, prompt injection, and quantum computing - that are specific to agentic payment systems.

5.2 **Know Your Agent and delegated authority**

- 5.2.1 The 'Know Your Agent' (KYA) concept was introduced as a framework to address agent identification - agents being neither natural persons nor legal entities within traditional KYC frameworks. Existing AML tools can trace wallet hops, but the proliferation of wallets (millions can be created in seconds) and the possible ephemerality of individual agents can render traditional identity verification structurally inadequate.
- 5.2.2 Two distinct dimensions of the identity problem were separated:

- (a). **AML and transaction monitoring.** The concern is transactional - following value flows and attributing transactions to responsible parties. A degree of pragmatism was noted: regulated frameworks with acknowledged gaps have historically been workable, and agentic payments may require a similar approach.
- (b). **Trust and identity.** The concern is relational - whether the merchant, PSP, or bank has sufficient confidence in the agent's identity and authority at the point of transaction. This is distinct from the AML problem and requires different technical solutions.

5.2.3 A W3C verifiable credentials framework was proposed: a human user issues a cryptographically signed credential to the agent, which presents it at the point of transaction, tying actions to the human principal without the agent requiring its own legal identity. This is compatible with ephemeral agents; the underlying technology exists but has not yet achieved widespread adoption.

5.2.4 A hierarchical structure - wallet issuer → individual → agent - was described as helpful for authentication and delegation verification. Fragmentation risk was noted: multiple competing KYA frameworks could replicate the heterogeneity of the current payments ecosystem, undermining the value of a centralised trust signal.

5.3 Liability allocation

5.3.1 In retail contexts with pre-defined mandates, the liability analysis maps clearly onto existing frameworks. Selecting the optimal product is not regulated activity (unless the product is a regulated financial instrument); initiating a payment instruction is, and falls within the PISP framework. The agent's operator must be regulated or operate through a regulated entity - no wholly new liability regime is required, only application of existing PISP rules to a new authorisation mechanism.

5.3.2 Corporate treasury use cases are more complex: dual signing authorities, board authorisations, and the scope of delegation to the agent require careful structuring. Multi-signature arrangements and key sharding - dividing signing authority across multiple corporate functions with a defined threshold required for authorisation - were raised as technically viable approaches. Algo trading regulation was noted as the closest existing regulatory precedent.

5.3.3 Merchants will pay for payment methods that demonstrably increase conversion, and an agent that shortens the payment journey has a clear economic case. This dynamic was identified as the most likely driver of widespread adoption: once merchants demand it, PSPs will follow.

5.4 **Agent legal personhood and agent-to-agent economies**

- 5.4.1 Under existing law, an AI agent (unless it is a legal person) cannot own assets, enter contracts, or bear liability in its own name. Where a blockchain records settlement between two wallet addresses, the law does not recognise it if neither belongs to a legal person - creating two parallel systems whose records may be irreconcilable in a dispute.
- 5.4.2 A speculative thread emerged around fully autonomous agent-to-agent economies. The legal framework does not accommodate these: ownership requires legal personhood, which requires a human or corporate controller. Such experiments are already occurring at the margins of the blockchain ecosystem, presenting a future challenge without a current legal solution.

5.5 **Fraud detection in an agentic world**

- 5.5.1 Existing fraud detection is calibrated to human behavioural patterns. In an agentic environment, bot activity is the legitimate activity, and models flagging automated behaviour as suspicious will generate false positives on normal authorised transactions. This is a foundational rather than incremental problem; no satisfactory solution currently exists, representing a genuine pre-competitive challenge for the industry.

5.6 **Prompt injection and systemic risk**

- 5.6.1 Prompt injection - embedding malicious instructions in content an agent reads (a webpage, email, or API response) to cause an unauthorised action such as a fraudulent payment - was raised as the paradigmatic agentic security risk. Unlike phishing, which targets individuals, it can exploit a shared vulnerability across thousands of agent instances simultaneously, at machine speed.
- 5.6.2 The likely market response is a certification regime - analogous to PCI-DSS - that agents must satisfy to achieve broad merchant acceptance. Uncertified agents may operate at the margins but will face limited acceptance within the regulated economy.

5.7 **Quantum risk to payment cryptography**

- 5.7.1 Google's quantum research has moved the plausible Q-day timeline to 2029 from the prior NIST reference window of 2030–2035, and several major infrastructure players have already moved their internal timelines accordingly. Quantum computing capability undergoes step-changes rather than gradual improvement, making 2029 a meaningful planning horizon.

5.7.2 Blockchain payment networks rely on ECDSA for wallet security. A quantum-capable adversary could derive private keys from public keys, gaining covert control of wallets at scale without detection. Migration is structurally difficult for decentralised networks - no central authority can mandate it. In the US, post-quantum cryptography compliance is now mandatory for regulated industries. AI compounds the threat by enabling attacks at unprecedented scale.

6. **SESSION 4: THE NEXT TWELVE MONTHS AND THE TRANSATLANTIC AGENDA**

6.1 This forward-looking session addressed what needs to happen in the next twelve months for agentic payments to move toward mainstream adoption, and examined where UK and US approaches diverge or converge.

6.2 **Currency infrastructure: can the UK avoid dollarisation?**

6.2.1 The session opened with a foundational question: can a UK agentic economy operate on GBP-pegged stablecoins, or does it default to tokenized deposits via the banking system or dollar-denominated coins? The discussion produced a nuanced picture.

- (a). For UK-facing services, sterling-denominated stablecoins are the natural preference, and the stablecoin market is likely to mirror the distribution of fiat currency use - retail service providers will want to transact in their customers' currency.
- (b). The earlier fear that the UK was so far behind that USD would become the de facto base, with an acceptance that technology had to be built around it, has receded somewhat but is not eliminated. If UK regulation does not deliver on its current trajectory, the risk re-emerges.
- (c). The GENIUS Act in the US was identified as a deliberate instrument to entrench dollar supremacy in the digital payments layer. The Eurozone and some African jurisdictions are actively pursuing protectionist digital currency regimes partly in response. The UK's position is less concerned: dollar dominance is less immediately threatening in the retail context, and the wholesale dimension has not yet generated an issue.
- (d). Currency sovereignty was not seen as the primary driver of UK policy in this space. Economic opportunity and job creation appear to take precedence over the international standing of sterling. Which currencies dominate will ultimately be determined by functionality, not geopolitical preference.
- (e). The Chinese model - combining tokenized bank deposits with, if approved, stablecoins - was noted as one alternative.
- (f). It was noted that the UK, and sterling, has gone down in terms of external use and attractiveness but there was optimism that with

the right infrastructure and marketing the UK could pick up some of this market share.

6.3 **Bank of England constraints on stablecoins**

- 6.3.1 The Bank of England's evolving approach to stablecoin regulation, specifically holding limits, was discussed as a material constraint on UK agentic payments.
- 6.3.2 The Bank is moving toward an issuance cap model for systemic stablecoins.
- 6.3.3 It was noted that UK lending markets are more closely tied to banking deposits than their US counterparts, such that competition between stablecoin holdings and bank deposits could pose sharper macroeconomic concerns than in the U.S.
- 6.3.4 Supporting stablecoins at scale requires a robust short-term money market infrastructure. UK money markets are more oriented toward the long term, and building the short-term infrastructure needed to back stablecoin issuance is a genuine policy gap. The government is consulting on this but has not yet resolved it.

6.4 **Are stablecoins necessary for agentic payments?**

- 6.4.1 A central debate in this session was whether stablecoins are a prerequisite for agentic payments, or whether existing infrastructure will suffice.
 - (a). Agentic programs appear easier to run using existing wallets, but stablecoins are not the only viable path. Stablecoins, as software-based instruments, interact more naturally with software-based agents, and have clear advantages in cross-border contexts - particularly in terms of speed and cost and when combined with the blockchain environment.
 - (b). An alternative path would be to make bank-backed money more programmable.
 - (c). In a decentralised environment, stablecoins become significantly more important. In an instant payment system, the comparative advantages are less clear.
 - (d). The decentralisation question itself has structural limits. Once a large, centralised actor enters the picture, the system is no longer truly decentralised. Efforts to reduce fees, and to reduce moving around assets, introduce reconciliation issues and raise questions about whether atomic settlement is genuinely achievable.

6.5 **A vocabulary gap**

6.5.1 A final observation concerned the merger of AI and token-based payment systems. Many actors in the AI community are unfamiliar with core concepts from the payments and tokenization world.

We note the proposed continuation of this discussion at DC Fintech Week in October, including the involvement of US regulators and US-based stakeholders.

Contacts

Clifford Chance



Simon Crown

Partner

Email: simon.crown@cliffordchance.com

Mobile: +44 79 0016 7265



Diego Ballon Ossio

Partner

Email: diego.ballonossio@cliffordchance.com

Mobile: +44 78 7039 8479



THE FINTECH FOUNDATION



Yesha Yadav

Milton R. Underwood Chair and Professor of Law
Associate Dean &
Robert Belton Director of Culture & Community
Faculty Co-Director, LL.M. Program
Board Member, Fintech Foundation

Email: yesha.yadav@law.vanderbilt.edu

Mobile: +1 615-322-2615

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 2026

Clifford Chance LLP is a limited liability partnership registered in England and Wales under no. OC323571. The firm's registered office and principal place of business is at 10 Upper Bank Street, London E14 5JJ. The firm uses 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** • Riyadh* • 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.