FCA CP25/40 Regulating cryptoasset activities | GBBC Response

About Global Blockchain Business Council (GBBC)

GBBC is the largest leading non-profit association for the blockchain, digital assets, and emerging technologies community. Founded in 2017 in Davos, Switzerland, GBBC comprises more than 500 institutional members and 251 Ambassadors across 119 jurisdictions and disciplines.

GBBC furthers adoption of digital technologies by engaging regulators, business leaders, and global changemakers to harness these transformative tools for more secure and functional societies.

GBBC industry verticals: Financial Services, Global Commerce/Supply Chains, and Commodities, underpinned by AI, digital identity, governance, hardware, infrastructure, policy, regulation, and security.

GBBC initiatives: BITA Standards Council (BITA), GBBC Digital Media, Global Standards Mapping Initiative (GSMI), International Journal of Blockchain Law (IJBL), InterWork Alliance (IWA), and U.S. Blockchain Coalition (USBC).

Executive summary

GBBC welcomes the FCA's continued engagement and the direction of travel in CP25/40 towards a workable UK regime that protects consumers and supports fair, orderly markets. We broadly support the FCA's objectives, particularly avoiding "post-box" models and ensuring UK supervision is meaningful.

Our central message is that several proposals in CP25/40 need tighter calibration so they do not unintentionally worsen outcomes for UK users by restricting access to global liquidity and weakening price discovery. Cryptoasset markets are structurally global; the UK framework will work best where it focuses on clear accountability and outcomes, rather than relying on assumptions drawn from geographically bounded markets.

We therefore recommend: (i) publishing auditable, outcomes-based criteria for "sufficient UK presence" so firms can plan predictable operating models; (ii) ensuring best execution expectations remain workable where UK-authorised venues or price sources do not exist at scale, including allowing reliable non-UK benchmarks under clear quality criteria; (iii) providing a clear menu of recognised conflict mitigants to support the move away from mandatory legal separation; and (iv) anchoring transparency and recordkeeping to consistent data standards (including a recognised asset identifier), with responsibilities allocated to the party best placed to hold the relevant data.

We also encourage the FCA to draw a clear line between custodial and non-custodial models for lending/borrowing and staking, so consumer protections apply where firms actually intermediate client assets and can practically meet the requirements.

Finally, we suggest early coordination with HMRC on potential interactions with the UK's Crypto-Asset Reporting Framework (CARF), as branch-plus-entity models could create duplicative reporting obligations for the same transaction if not addressed upfront.

GBBC has focused its response on the consultation questions most relevant to our remit and to areas where we can provide the greatest value based on our members' expertise and practical experience.

Consultation Questions

Question 1: Do you agree with our proposals on location, incorporation and authorisation of UK CATPs? If not, please explain why not?

We support the FCA's intent to prevent hollow "post-box" structures and ensure the regulator has genuine supervisory reach into UK-facing operations. Where we push back is on how the current framing could impose unnecessary governance ambiguity on firms running legitimate, integrated global platforms, not because they want to evade UK oversight, but because the architecture of cryptoasset markets is fundamentally different from the domestic equity or derivatives markets the existing framework was built around.

This matters practically. For most tokens, the deepest liquidity and most efficient price discovery happen on consolidated global order books. If the rules, whether by design or by implication, push firms toward a subsidiary model with a distinct UK order book, the result won't be better protection for UK users; it will be fragmented liquidity, wider spreads, weaker price discovery, reduced asset availability, and a quiet migration of activity offshore. These aren't hypothetical risks, they're the predictable structural consequences of imposing local market-microstructure requirements on a global, 24/7 asset class.

That said, the FCA's underlying objective is legitimate, and firms in this space have an interest in a clear, supervised perimeter. The real question is how to deliver meaningful UK control and accountability without compelling local liquidity replication as a side effect. Three things would meaningfully improve the current framework.

First, the FCA should articulate clear, minimum "UK nexus" criteria that are outcomes-based but auditable covering SMF accountability, UK-resident governance and escalation, UK control rights over key risk decisions (market integrity controls, incident response), access to relevant systems and data, and a demonstrable ability to direct incident management in real time. Framing this around supervisory control rather than structural form would allow the FCA to achieve its objectives without inadvertently prescribing operating-model outcomes it didn't intend.

Second, the CP signals a degree of firm-by-firm flexibility on how the location test is applied. Flexibility is useful, but without a published baseline it effectively means the compliance answer is negotiated each time, which creates planning risk for firms and inconsistency in outcomes across the market. A standardised accountability map setting out what must be UK-resident versus what may remain within a global operating hub would reduce

supervisory friction, give firms something to plan against, and produce more consistent regulatory outcomes.

Third, where the primary policy concern is UK retail access and conduct, it's worth the FCA considering whether an authorised intermediating broker can serve as the supervised interface between UK consumers and global liquidity, rather than requiring the trading venue itself to fragment its liquidity pool in every case. This preserves global price formation while keeping the UK conduct and supervisory perimeter intact via the regulated intermediary. It's not the right answer in every structure, but it deserves to remain a live option.

One further issue that we don't think has received sufficient attention is the interaction between the proposed branch-plus-entity model and the UK's CARF implementation. From 1 January 2026, reporting crypto-asset service providers carry obligations to report on both resident and non-resident users. The operating model the CP contemplates where a global provider runs via both a UK incorporated entity and a separate UK branch raises a straightforward but underexplored risk: the same exchange transaction becoming reportable twice in the UK, once through the entity and once through the branch, on behalf of the same user. CARF includes switch-off mechanisms where reporting occurs in a Partner Jurisdiction, but there is no equivalent provision addressing duplicative reporting within the same jurisdiction across a branch-and-entity structure of the same group.

We believe that this is a direct consequence of the structural model under discussion, and it creates a compliance burden without improving supervisory outcomes. We'd encourage the FCA to work through this with HMRC before finalising the framework, with the goal of establishing clear guidance on how reporting responsibility should be allocated across branch/entity structures, or a mechanism that ensures a single UK reporting obligation per transaction regardless of structural form.

Several proposals in the CP interact in ways that could unintentionally close off access to global liquidity tools. In particular, the CP seeks to preserve global order-book efficiencies while also constraining where UK retail orders can be executed and limiting the extent to which UK-authorised intermediaries can rely on offshore venues. Taken together, these design choices risk creating a de facto “UK-only” market structure that is unlikely to replicate global price formation and could worsen outcomes for UK users.

Question 2: Do you agree with our proposals on UK CATP access and operation requirements? If not, please explain why not?

We support the FCA's objective of ensuring that UK-facing cryptoasset trading occurs within an accountable supervisory perimeter. The question is whether the access and operating model proposals can achieve that without inadvertently severing UK participants from the global liquidity and price formation that actually determines execution quality.

This matters structurally. Cryptoasset markets are global by design. For most tokens, consolidated liquidity and efficient price formation happen on global order books, not on fragmented national pools. A framework that constrains access to offshore venues, whether directly or through restrictions on UK intermediaries' ability to route to foreign platforms, doesn't just impose compliance costs. It closes off the liquidity sources that deliver good outcomes for UK users in the first place.

Our concern is that design choices across different parts of the CP, when combined, could leave firms unable to access global liquidity in practice. That's not a theoretical risk. If UK retail orders must execute on UK-authorised venues, and those venues cannot themselves access or replicate the depth and efficiency of global order books, the result is predictable: wider spreads, weaker price discovery, reduced asset availability, and incentives for sophisticated participants to move activity offshore while retail users are left with a lower-quality product.

In that context, we encourage the FCA to consider whether the most effective and proportionate point of supervisory control is the UK-regulated intermediary interface, where UK consumers interact with markets, rather than requiring the trading venue itself to replicate liquidity locally. A model in which UK-authorised brokers or intermediaries provide the controlled interface to global venues can maintain strong consumer protections and FCA oversight while avoiding structural fragmentation of liquidity.

We also recommend the FCA provides greater clarity on perimeter interactions. Uncertainty about whether non-UK entities can safely access UK venues without inadvertently entering the UK perimeter creates unnecessary legal and commercial friction.

Additionally, it is important that access and operational requirements explicitly support market integrity outcomes in a market where abusive conduct can occur both off-chain and on-chain. In practice, this means CATPs should maintain surveillance capabilities that monitor relevant on-chain activity on a risk-based basis, alongside traditional off-chain monitoring, to support detection and prevention of market abuse. Where algorithmic or automated trading represents material activity on a platform, there should be clear expectations that platforms monitor algorithmic trading for compliance with venue rules and market abuse obligations, and provide appropriately high-level disclosure on the significance and nature of that trading.

Question 3: Do you agree with our proposals on additional rules to protect UK retail customers? If not, please explain why not?

We support robust protections for UK retail customers, but those protections need to be designed around market structure realities rather than imposed in ways that unintentionally degrade the outcomes they're meant to improve.

The critical design question is where retail protection is anchored. The most effective point of control is where UK consumers are actually served, through UK-authorised intermediaries subject to clear conduct standards and FCA oversight. That's where accountability can be enforced directly, and where the FCA's supervisory reach is strongest. Structural requirements that go beyond that point and attempt to fragment liquidity or constrain access to global price formation don't add meaningfully to consumer protection. What they do add is execution cost: wider spreads, weaker price discovery, and reduced access to assets. For retail users, that's not protection. It's a worse product.

The risk here is that the framework forces retail execution into UK-only liquidity pools, whether directly through venue restrictions or indirectly through constraints on how intermediaries can route orders. If that happens, the policy will have succeeded in keeping activity onshore while simultaneously making that activity less beneficial to the consumers it was designed to protect. That's an avoidable outcome, but avoiding it requires calibrating the protections carefully.

We therefore recommend the FCA design retail protections to secure clear UK accountability and conduct standards at the intermediary level, while preserving firms' ability to access global liquidity and demonstrate best execution using high-integrity sources where UK liquidity is insufficient. That balance is achievable, but it requires the FCA to distinguish between the supervisory perimeter it needs to control and the market structure it should allow firms to access within that perimeter.

Question 4: Do you agree with our proposals to manage conflicts of interest and related risks? If not, please explain why not?

We support an outcomes-focused approach to conflict management, and in particular the shift away from treating legal separation as the default solution. One of the things the FCA could do is publish a recognised menu of mitigants that firms can combine proportionately depending on their business model: independent surveillance capability with clear reporting

lines, information barriers between venue operations and proprietary trading, governance committees with genuine escalation and veto rights, restrictions on affiliated liquidity sourcing where it undermines fairness, and audit trails demonstrating independence in decision-making and incident handling. Without clear signals from the FCA about which combinations are sufficient, the regime risks becoming one where compliance is only confirmed retrospectively.

We also recommend the FCA explicitly recognises that effective conflict mitigation in crypto markets often depends on surveillance that can join up on-chain and off-chain signals, particularly where conflicts can manifest through trading behaviour. In addition, we reiterate the value of the FCA providing non-exhaustive illustrative examples of controls that may be appropriate for different conflict types - affiliate trading, market making, matched principal activity - to support consistent implementation while preserving proportionality.

There is also a structural interaction worth noting. If the overall execution architecture constrains access to global liquidity, firms may end up in models that generate more conflicts risk rather than less. A venue that cannot access global price formation efficiently has stronger incentives to rely on affiliated liquidity. The conflicts management framework should be designed with this dynamic in mind.

Question 6: Is any further guidance on best execution required? If so, what additional guidance can we provide to clarify the scope of and expectations around best execution?

Further guidance is necessary in three areas. Firstly, the relationship between the rules and how evidential compliance actually works over time. Firms need to understand what good looks like for order execution policy governance, how periodic testing should be structured and documented, and how exception handling feeds back into policy review. The FCA should clarify that that best execution is not assessed as a rigid trade-by-trade mechanical benchmark. That framing, if left ambiguous, distorts firm behaviour in ways that don't serve clients: it incentivises gaming of narrow metrics rather than genuine focus on execution quality across the range of factors that actually matter.

The second area where guidance would add real value is alignment with how best execution is already understood for comparable market activity. Where a firm can demonstrate through data that a single execution venue consistently delivers best outcomes across the relevant execution factors, that should be a sufficient basis for compliance without requiring artificial diversification of venues or a permanent obligation to re-litigate the analysis. Bringing the

cryptoasset framework into line with how this works in adjacent markets would reduce unnecessary divergence and leverage the substantial body of practice that already exists.

The third issue is that best execution in cryptoassets often depends on accessing price formation that happens outside UK-authorised venues. A framework that, in practice, limits firms to UK-authorised price sources will produce worse execution outcomes, not better ones and will be difficult to apply honestly by firms that know the real price is being discovered elsewhere. The FCA should establish a clear framework under which high-integrity global price sources can be used to demonstrate best execution, subject to defined criteria around source quality, independence, and auditability.

Question 7: Do you agree with our proposed guidance (including the exemptions proposed) to check at least 3 reliable price sources from UK-authorised execution venues, such as a CATP or principal dealer (if available)? If not, please explain why not?

We agree with the underlying principle that firms should survey the wider market rather than defaulting to a single venue without proper justification. The difficulty arises from the assumption that three UK-authorised execution venues providing genuinely independent price signals will be available for the assets firms need to trade. For a significant proportion of tokens, they will not be. A numeric expectation under those conditions either becomes impossible to satisfy or distorts order routing toward a limited set of venues not because they deliver best execution, but because compliance requires it.

There is also a question about how this guidance will operate in practice. When guidance includes exemptions, it creates uncertainty about whether firms can genuinely treat the expectation as non-binding or whether it will effectively become a supervisory requirement. This matters particularly for firms operating across both cryptoassets and derivatives, who need execution governance that works coherently across asset classes. A crypto-specific numeric rule that doesn't align with how best execution operates for derivatives creates operational complexity without improving client outcomes.

A further consideration is international consistency. Restricting benchmark or reference prices by UK authorisation status may push firms toward less representative pricing inputs even where higher-integrity, deeper-liquidity reference venues exist elsewhere. A more coherent approach would establish criteria for source reliability and integrity - governance, controls, surveillance, data quality - irrespective of location. Best execution should remain

anchored to market reality while remaining fully auditable, rather than being artificially constrained by authorisation geography.

What would work better is framing this as a proportionate, policy-level expectation rather than a numeric rule. Firms should demonstrate through their order execution policy, and through how that policy is governed and reviewed over time, that they systematically assess execution quality across available venues and do not anchor to a single venue without justification. This standard produces good outcomes without creating implementation problems for tokens where UK-authorised venues are scarce.

Where UK-authorised sources are genuinely insufficient - either too few in number or not reflective of where price formation actually occurs - firms should be permitted to use non-UK sources that meet defined criteria around venue integrity, surveillance and controls, data reliability, and transparency. The standard should be source quality, not source location. Governance and documentation requirements ensure this remains auditable without forcing firms to disregard where prices are genuinely being discovered. For assets where the deepest liquidity and most efficient price formation occur on global platforms, best execution for UK clients depends on accessing those sources, not relying on UK-authorised alternatives that may be inadequate.

Question 11: Given the overall location policy established by the amendments to section 418 of FSMA set out in the Cryptoasset Regulations, do you agree with our proposed execution venue requirement? If not, please explain why not?

We recognise the policy objective here: ensuring UK retail orders aren't routed offshore without appropriate oversight, and anchoring execution within a supervisory perimeter the FCA can actually reach. The concern is that an "execute only on UK-authorised execution venues" requirement, applied as a hard rule to retail and elective professional flows, risks causing exactly the harm it is designed to prevent: UK consumers ending up with worse execution outcomes, not because firms are evading supervision, but because the liquidity and price formation they need is global and cannot simply be relocated by regulatory instruction.

The CP itself acknowledges this tension - the references to cliff-edge risks for existing holders, the need for targeted exemptions, and the recognition that transitional arrangements will be necessary are all signals that the FCA understands the market structure problem. Our concern is that those acknowledgements haven't yet translated into a framework that resolves it. Carve-outs and transitional relief applied to an otherwise rigid requirement tend to produce

compliance complexity and supervisory uncertainty rather than coherent, durable operating models.

The more productive framing is to start from the FCA's actual policy goals - consumer protection and effective supervision - and ask what structural conditions are necessary and sufficient to deliver them. On that analysis, the critical variable is whether there is a UK-supervised entity standing in the chain with accountability for conduct, client outcomes, and regulatory compliance, not whether the venue itself executes against a locally replicated order book. Where a UK-authorised intermediary sits as the controlled interface between the UK client and the execution venue, with appropriate conduct obligations, oversight of execution quality, and the FCA's ability to reach it directly, the supervisory objective is met. Requiring the venue to onshore its liquidity in addition to that adds structural cost and execution quality degradation without adding meaningfully to consumer protection.

Where the FCA does retain an execution venue requirement, we recommend it be complemented by a clear and proportionate recognition pathway for overseas venues subject to comparable regulatory and supervisory outcomes. This would help preserve execution quality and reduce liquidity fragmentation while remaining consistent with the FCA's location and enforcement objectives, particularly where UK firms can demonstrate robust governance over venue selection, execution quality monitoring, and effective client protections.

The FCA should also clarify how global group structures can support UK operations without falling foul of the policy intent, including safe operating model guidance for firms running integrated platforms where UK-facing activity is a component of a larger global operation. Without that clarity, firms face structural choices under uncertainty, and the market design outcomes - fragmented liquidity, reduced asset availability, incentives to restructure in ways that move activity rather than bring it under oversight - will be worse for UK clients than the status quo the regime is trying to improve on.

Question 13: Do you agree with our proposed approach to addressing conflicts of interest during order execution when a firm is engaged in proprietary trading? If not, please explain why not?

The proposed conflicts approach is sound in principle, and we support the FCA's focus on ensuring fair treatment of clients where firms engage in proprietary trading. Our observation is that conflicts mitigation works best when it's designed in light of the broader execution and

liquidity architecture, and there may be ways to align these two parts of the framework more closely.

One area worth considering is how liquidity access interacts with conflicts management. Where firms have ready access to deep, external liquidity sources, it becomes easier to demonstrate fair treatment across participants and to manage conflicts through structural separation and independent price validation. If the framework were to constrain access to global liquidity, whether by design or as an unintended consequence, that could make conflicts management more challenging rather than less so. Venues with fewer external options may face stronger incentives to rely on affiliated sources, which is precisely the dynamic the conflicts rules are meant to address. We'd encourage the FCA to think through this interaction carefully.

There may also be merit in considering how different structural models affect conflicts risk. In some cases, a model where global venues continue to operate globally, with UK retail participation intermediated through UK-authorised brokers, could provide a cleaner separation of responsibilities. The venue focuses on neutral price formation across its participant base, while the intermediary under FCA supervision manages the client relationship and execution quality. This isn't the right answer for every business model, but it may offer a route to reducing conflicts structurally in certain contexts.

On the requirements themselves, greater clarity on recognised mitigants and evidential expectations would help firms design robust, auditable controls from the outset. The underlying components are well understood - surveillance independence, information barriers, governance oversight, audit trails - but guidance on which combinations meet the FCA's expectations and what standard of evidence is required would improve predictability and consistency across the market.

Question 17: Do you agree with our proposed pre-and post-trade transparency requirements for UK CATP operators and principal dealers? If not, please explain why not?

We support the direction of travel on pre- and post-trade transparency. Better price discovery is a genuine public good in these markets, and a framework that encourages consistent reporting practices is worth building. The first concern is one of signal quality rather than compliance cost. Cryptoasset price formation is global. Where the significant majority of volume and liquidity in a given token trades on venues outside the UK regime, transparency requirements applied only to UK activity will produce a partial and potentially misleading

picture of market conditions. Transparency data that captures a thin slice of real activity can actively distort it by creating reference points that don't reflect where the market actually clears. Cryptoasset price formation is global, so UK-only transparency will often be a thin and potentially misleading slice of the true market. The FCA should therefore (i) make clear that the primary purpose is supervisory oversight of UK-authorised activity (not creation of a UK benchmark), (ii) calibrate pre-trade transparency by liquidity profile with clear waivers/deferrals and permission to use consolidated reference pricing where UK liquidity is immaterial, and (iii) provide specific guidance for global order-book models so transparency does not inadvertently force a UK "data silo" that neither improves price discovery nor execution outcomes.

The second issue is more technical, we agree with the requirement for a "unique and unambiguous" cryptoasset identifier. However without a consistent identifier, transparency data can't be aggregated, compared, or used programmatically in any meaningful way. The identifier needs to be anchored to a recognised standard. This mitigates firms making e different implementation choices resulting in fragmented data. There are existing frameworks the FCA could reference, ISIN extension work for cryptoassets, DTIF, or other emerging standards and the FCA should specify clearly which it expects firms to use, or set out the criteria a qualifying identifier framework must meet. Leaving this open creates interoperability risk that undermines the usability of the transparency data the regime is designed to generate. As firms will need to build infrastructure around whatever standard is specified, early clarity is necessary.

Question 18: Do you agree with our proposed methodology for determining the pre-trade transparency threshold? If not, please explain why not? What other methodology do you suggest?

We agree a threshold is necessary - applying pre-trade transparency requirements uniformly regardless of scale would impose disproportionate cost on firms whose contribution to price formation and liquidity is modest. The policy rationale for calibrating scope is sound. The concern is with whether a single £10m annual revenue figure applied at the entity level is the right instrument to achieve that calibration.

The mismatch problem is straightforward. A diversified firm with substantial revenue from activities unrelated to cryptoasset dealing or CATP operations could clear the threshold comfortably while running in-scope crypto activity that is small in absolute terms and immaterial to market-wide price formation. Conversely, a firm whose business is concentrated in cryptoassets but whose revenue profile is more modest might fall below the

threshold despite being a meaningful liquidity contributor in specific tokens or market segments. In both cases, the threshold produces outcomes that don't track the policy objective to capture firms whose transparency obligations would generate data of genuine market value, not simply to capture large firms by revenue regardless of what they do.

The FCA itself acknowledges alternatives: transaction volume, trade count, or a measure of liquidity contribution are all closer proxies for what actually matters than total entity revenue. In our view a composite approach has merit: a revenue floor to exclude genuinely small operations, combined with an activity measure to ensure the threshold tracks in-scope business rather than group size. That combination preserves simplicity at the lower end while avoiding the misclassification problem for diversified firms. We also suggest a building in a formal review, either a sunset clause or a committed calibration exercise, once the regime has generated a first wave of market data. The right threshold for a nascent regulatory framework is unlikely to be the right threshold three years in, and locking calibration choices in without a review mechanism stores up problems as market structure evolves.

Question 19: Do you agree with our proposals for transaction recording and client reporting requirements for UK CATP operators and intermediaries? If not, please explain why not?

We support robust transaction recording and client reporting obligations. The ability to reconstruct the full lifecycle of a retail order, and to do so efficiently under supervisory scrutiny or in the context of a dispute, is essential. The implementation question is how to allocate those obligations sensibly across execution chains that involve multiple parties playing different roles, some of which will have far better visibility into certain parts of the record than others.

We consider the principal risk here not under-recording, but misallocated recording: requirements that land on parties that do not hold the relevant data, or that result in multiple parties independently maintaining overlapping records of the same events in inconsistent formats. Both outcomes create cost without improving supervisory access, and the inconsistency problem can make reconstruction harder rather than easier when records from different parts of the chain need to be reconciled.

Responsibility should map to visibility. In a CATP-executed model, the CATP holds the most complete picture of execution and is the natural home for execution-side records. In a principal dealer model, the dealer's records of the transaction it internalised are the primary source. In a routed or intermediated structure, the split of responsibility between the

intermediary and the downstream venue needs to be explicit: which party is accountable for which fields, and how the records are expected to be linked across the chain.

Where transactions touch public networks, firms should also be able to link off-chain execution records to relevant on-chain artefacts in a way that supports auditability and reconstruction. In practice, this means retaining, where available and relevant, transaction hashes, wallet addresses or other network identifiers, and network/block timestamps or settlement references, mapped back to the client order identifier and the execution record. This should not be framed as a requirement to "mirror the blockchain" off-chain; rather, it is about maintaining reliable join keys between internal order/execution systems and the underlying network events so that an order can be traced end-to-end without ambiguity.

On content, requirements should focus on the fields that are hardest to reconstruct after the fact and most critical for end-to-end traceability: consistent and synchronised timestamps across the execution chain; execution venue identifiers that are unambiguous and stable; and asset identifiers that tie back to a single FCA-endorsed standard (and are consistent across transparency, disclosures, and recordkeeping obligations), so that the same cryptoasset is not represented differently across datasets. These are the fields where gaps or inconsistencies cause real supervisory and operational problems; standardising them delivers significant value.

We also recommend that recordkeeping expectations explicitly capture responsibility and decision-making within the execution process, including where decision points are automated. Where smart order routing or algorithmic logic determines venue selection, order handling, or execution sequencing, firms should retain sufficient records to evidence the basis on which decisions were made, the parameters in force at the time, and the relevant controls and governance over that logic. This supports both dispute resolution and supervisory investigations where the question is not only what happened, but why it happened.

Standardised data formats would multiply these benefits. Firms that can satisfy recording requirements using common schemas, and share or compile records across platforms without bespoke translation work, face materially lower compliance cost, and the FCA gets data that can actually be aggregated and interrogated rather than a patchwork of proprietary formats. Early clarity from the FCA on expected standards would have significant practical value. In particular, we encourage the FCA to provide a non-exhaustive set of illustrative "minimum viable" record fields for common operating models (CATP-executed, principal dealing, routed execution), and to clarify which party is expected to be the system of record for those fields.

Question 21: Do you agree with our proposal to prohibit the use of proprietary tokens for L&B as outlined above? If not, please explain why not?

We understand the consumer protection rationale here, and where proprietary tokens are being used in custodial lending and borrowing models to create or obscure conflicts of interest or to engineer incentive structures that benefit the platform at the expense of clients: a prohibition is a reasonable policy response. That case is clearest where the firm controls both the token and the custody of client assets, and where clients may not fully appreciate how the token's use shapes the economics of the arrangement they're entering into.

The concern is about scope. The prohibition as currently framed risks capturing models where a native token isn't functioning as a conflicted incentive instrument at all; where it is instead structurally integral to how a non-custodial protocol operates and governs itself. In those contexts, the token isn't something layered on top of a lending arrangement to benefit a controlling entity; it is part of the mechanism by which the protocol functions, allocates risk, and makes governance decisions. Applying a prohibition designed for custodial conflicts to that kind of structure would be a category error, one with real consequences, since it would effectively make compliant operation of certain non-custodial models impossible without restructuring that serves no consumer protection purpose.

In our view, the FCA should explicitly limit the prohibition to custodial lending and borrowing models, or at minimum draw a clear and operationally usable distinction between custodial and non-custodial structures in how the rule is framed and applied. If the prohibition is intended to address the conflict that arises when a firm controls both the token and the client's assets, then draw the perimeter accordingly. If there are specific non-custodial arrangements that also raise material concerns, those should be identified and addressed on their own terms rather than caught incidentally by a broadly-framed rule.

There is also a coherence point across the wider regime. The question of when a token or protocol is subject to a firm's "control" and therefore within the scope of rules designed for custodial or intermediated models, is one that will arise in staking and DeFi-adjacent contexts as well as in lending and borrowing. The FCA has signalled that guidance on "control" is forthcoming, and the final approach to this prohibition should be aligned with that guidance rather than allowed to develop independently. Inconsistent scope across L&B, staking, and DeFi-like models creates interpretive risk for firms trying to structure compliant operations across multiple product lines, and resolving those inconsistencies after rules are in force is substantially harder than building coherence in from the outset.

Question 22: Do you agree with our proposed record-keeping requirements on regulated L&B firms? If not, please explain why not?

The record-keeping requirements as proposed are well-suited to custodial lending and borrowing arrangements, where firms hold client assets and maintain direct client relationships.

Our concern relates to scope and implementation across different structural models. The requirements as drafted appear to assume an intermediated arrangement - a firm holding client assets, maintaining account relationships, collecting KYC information, and establishing contractual terms. That model corresponds to custodial operations and the obligations fit naturally. The question is how these requirements apply to non-custodial protocols, where no central intermediary holds assets or operates the kind of client relationship the record-keeping regime contemplates.

Extending custodial-style record-keeping to non-custodial structures without modification does not address a supervisory gap. It establishes obligations that lack a clear implementation path because the intermediated relationship the requirements were designed around is absent. The entity expected to maintain records may not exist, or may not hold the data the regime assumes will be available. This creates a compliance problem that cannot be resolved through better systems or processes - the underlying structure does not support the requirements as written.

We recommend the FCA provide explicit guidance distinguishing how record-keeping obligations apply to custodial versus non-custodial arrangements. For custodial operations, the current framework is appropriate and proportionate. For non-custodial structures, the FCA should clarify either that certain requirements do not apply, or how they should be interpreted given the absence of a traditional intermediary holding client assets.

Question 25: Do you agree with our proposal that regulated staking firms must provide retail clients with information on the firm and its staking service, and provide the key terms of agreement in relation to those services and obtain retail clients' express prior consent in relation to those terms each time cryptoassets are staked, as outlined in paragraphs 6.14-6.19? If not, please explain why not?

It is essential that the regulatory framework draws a clear and explicit distinction between custodial and non-custodial staking, as these models involve materially different levels of intermediation and risk. Without precise scoping, there is a real risk that obligations designed

for firms that safeguard client assets are inadvertently applied to arrangements where providers neither hold private keys nor control user funds. Consumer-protection requirements such as disclosures, clear contractual terms, express prior consent and record-keeping are appropriate and proportionate in custodial staking, particularly for retail clients, because the firm intermediates rewards and exposes users to operational and counterparty risk. Those same obligations, however, are not suitable for non-custodial or protocol-level participation.

A functional differentiation between staking models is therefore critical. Delegated staking, pooled custodial staking, liquid staking and direct protocol-level staking represent distinct technical and legal activities with materially different risk profiles. Regulatory obligations should arise only where asset custody or safeguarding is actually present, rather than treating all forms of staking under a single conceptual label that conflates infrastructure participation with financial intermediation. In delegated or protocol-level staking arrangements, validators typically operate nodes without taking possession of client assets or private keys, meaning ownership and control remain with the user at all times. Where firms merely provide technical interfaces or communication infrastructure that enable interaction with a protocol, the activity is in substance facilitative rather than custodial and should fall outside the scope of regulated staking.

Aligning the UK approach with international best practices would materially improve clarity and consistency. Recent guidance in the United States and the European Union both anchor regulatory obligations to the presence of custody and managerial discretion rather than to participation in consensus mechanisms themselves. These frameworks distinguish protocol-native or self-custodial staking from intermediary-based services, ensuring that licensing and conduct requirements attach only where firms exercise control over client assets or materially shape economic outcomes. Adopting a similarly custody-linked, function-based distinction would enhance legal certainty, maintain proportionality, and keep regulatory intervention focused on the activities that genuinely transfer risk to consumers.

Question 27: Do you agree with our proposed record-keeping requirements on regulated staking firms? If not, please explain why not?

Record-keeping requirements are appropriate for custodial staking services where firms safeguard client assets or exercise discretion over staking operations. In those arrangements, comprehensive records support consumer protection and supervisory oversight.

The scoping issue is the same as for lending and borrowing. Where staking is non-custodial - protocol-level or smart-contract participation without the provider holding client assets or private keys - the record-keeping obligations assume a firm-managed relationship that may not exist. The relevant records may be on-chain rather than maintained by any particular

entity, and the provider may have no practical means to collect the off-chain data the proposals contemplate.

The FCA should link these requirements explicitly to custodial staking models. Obligations should attach where custody, control, and a client relationship exist. They should not be extended to non-custodial participation that does not involve safeguarding client assets, where the underlying structure cannot support the requirements as drafted.

Question 28: Do you agree with our proposal to apply rules and guidance in chapters 2-6 and guidance to firms engaging in DeFi where there is a clear controlling person(s) carrying on one or more of the new cryptoasset activities? If not, please explain why not?

We agree with the FCA's proposed approach as a starting point that extending regulatory obligations to DeFi where there is a clear controlling person or coordinated group is both principled and practically necessary. A framework that can be circumvented simply by labelling an arrangement as decentralised would be ineffective, and the FCA is right not to let form determine substance here. Building the control test in a way that is rigorous enough to catch genuine control without capturing infrastructure that is, in any meaningful sense, nobody's to control is a complex task.

We recommend the FCA build the framework around five principles:

1. **Common Control Test:** Regulatory obligations should hinge on whether a person or coordinated group can exercise material and unilateral influence over a protocol's operation or governance. Objective criteria, such as upgrade rights, voting power concentration, permissionlessness, and practical ability to direct outcomes, should be used, alongside clear safe harbours for networks that are not under common control.
2. **Public Permissionless Infrastructure Exemption:** Genuinely open, permissionless infrastructure that allows anyone to validate or participate without approval, and that does not involve custody or discretion, should fall outside the scope of regulated activities. Autonomous smart-contract protocols operating purely as neutral infrastructure should not be treated as financial intermediaries.
3. **No Substantial Dependence on Any Person or Group:** A key indicator of decentralisation is whether the system can continue functioning as intended without the ongoing discretionary involvement of any identifiable actor. If the protocol remains operational even when original developers or operators disengage, this strongly suggests absence of control.

4. **Limits on Unilateral Economic Control:** Clear quantitative thresholds should be used to assess dominance over token supply or governance power. If a single person or coordinated group controls a large share of token supply (20%), this is a strong signal of potential control.
5. **Decentralisation Spectrum:** Control and decentralisation are not binary but evolve over time. Protocols may decentralise as governance disperses and code becomes immutable, or re-centralise through token concentration or governance capture. Regulatory treatment should therefore be dynamic and calibrated to the current, demonstrable level of control rather than fixed classifications.

The most important structural implication of principle five, and one the framework needs to be designed around from the outset, is that a one-time classification will systematically misclassify protocols at both ends of the decentralisation journey. There must be clear mechanisms for reclassification in both directions, and reasonable expectations on firms and developers about when they must re-engage with the supervisory process as their protocol's governance structure evolves.