




WARWICK BUSINESS SCHOOL
THE UNIVERSITY OF WARWICK

Gillmore Centre for
Financial Technology

Response to BoE 2025 consultation paper:

Proposed regulatory regime for sterling- denominated systemic stablecoins



The Gillmore Centre for Financial Technology at Warwick Business School, University of Warwick, is a multidisciplinary research hub focused on the transformative impacts of emerging technologies, including artificial intelligence, machine learning, blockchain, mobile payments, cryptocurrencies, and crowdfunding platforms.

We welcome the Bank of England's 2025 consultation paper on its proposed regulatory regime for sterling-denominated systemic stablecoins and appreciate the opportunity to respond. This document aims to engage with several issues, ideas, and questions posed in the consultation. Our response is not exhaustive but explores diverse arguments and perspectives, while highlighting critical questions and initial thoughts on potential answers for the Bank's consideration. We hope it can be a conversation starter.

We would be pleased to connect with anybody interested in discussing or collaborating on these or other relevant issues.

Yours sincerely
Gillmore Centre for Financial Technology

Contributing authors:

- Basil Sansom, PhD – Research Fellow, WBS Gillmore Centre Financial Technology
- Prof. David Skeie – Professor of Finance and a research fellow at the Gillmore Centre for Financial Technology at WBS
- Natasha de Teran – Honorary Research Fellow, WBS Gillmore Centre Financial Technology
- Martin Walker – Honorary Research Fellow, WBS Gillmore Centre Financial Technology
- George Pantelopoulos, PhD – Lecturer in Economics, the University of Newcastle Australia
- Rhomaios Ram, Honorary Research Fellow, WBS Gillmore Centre Financial Technology
- Prof. Dalvinder Singh – Professor of Law, University of Warwick Law School
- Olga Klein, Associate Professor of Finance and a research fellow at the Gillmore Centre for Financial Technology at WBS
- Ganesh Viswanath-Natraj, PhD – Assistant Professor of Finance and a research fellow at the Gillmore Centre for Financial Technology at WBS
- Ka Kei Chan, Associate Professor of Finance and a research fellow at the Gillmore Centre for Financial Technology at WBS.

Corresponding author: bazil.sansom@wbs.ac.uk
Gillmore Centre: gillmorecentre@wbs.ac.uk

Disclaimer

This report is the product of the Gillmore Centre's working group (as listed) on innovation in money and payments and reflects broad agreement among its participants. This does not imply agreement with every specific observation or nuance. Members participated in their personal capacity, and their participation does not imply the support or agreement of their respective public or private institutions.

Contents

1. Summary of response.....	4
2. CP Questions responses	9
2.1 Question 1: Backing assets.....	9
2.1.1 Why allowing securities investment is not just a business-model tweak.....	9
2.1.2 Benchmark: what 100 % BoE deposits buys (and what it shifts onto banks).....	9
2.1.3 What the hybrid 40/60 backing changes	9
2.1.4 Achieving “£1 = £1 everywhere”: interoperability at the coin↔deposit boundary.....	11
2.1.5 Why repo access is needed and reverse repo should be considered.....	11
2.1.6 Liquidity backstop essential for payments grade par under hybrid backing.....	12
2.1.7 Macro and monetary footprint (why the sign is ambiguous and driven by 60/40).....	13
2.1.8 Non-remuneration anchors payments use; remuneration opens portfolio channel.....	14
2.2 Question 2: step-up regime	14
2.3 Question 3-6: Risk-based capital and reserve requirements.....	15
2.3.1 “Trust” model good, but semantic risk can pull understanding toward an investment-share mental model (“owners not creditors”)	15
2.4 Question 7: Other comments or suggestions on Section 2.2	16
2.5 Question 8-10: Holding limits and the risks they address	16
2.5.1 Operational challenges (and why they’re not just technical issues).....	17
2.5.2 Identity and aggregation: the hard problem is not the limit, it is “who is counted”	17
2.5.3 Where enforcement sits: on-chain strength versus governance and resilience burdens.....	17
2.5.4 Failure modes and user experience: limits without “waterfall” are hard constraints.....	17
2.5.5 How holding limits can turn “money” into “rails”	17
2.5.6 Usability depends on whether stablecoin is “money” or a corridor-dependent “rail”	18
2.5.7 Individuals: what remains feasible, and what really drives usability	18
2.5.8 Businesses: the decisive constraint, and why exemptions are not edge cases.....	18
2.5.9 What uses would be structurally discouraged or impossible under binding limits.....	19
2.5.10 How holding limits shape competition and gatekeeping.....	19
2.5.11 Tools other than holding limits (and how to position them).....	19
2.5.12 Clarifying what holding limits are trying to achieve in this CP	19
2.5.13 Why stock caps are a blunt instrument for run dynamics	20
2.5.14 Reduce the store-of-value pull (preferred first line of defence)	20
2.5.15 Make the corridor payments-grade by construction (avoid boundary fragility)	20
2.5.16 Flow tools at the boundary (if throttling is contemplated at all)	20
2.5.17 Perimeter and structural tools (enforceability, interoperability, and transition discipline)	20
2.5.18 Note on other limit types considered in the CP	20
2.6 Question 11: Direct access.....	21
2.7 Question 12-13: Backing asset safeguarding and proposed trust arrangements.....	21
2.8 Question 14-16: Public permissionless ledgers.....	23
2.8.1 Collapsed bearer model versus operator-recognised two-layer model.....	23
2.8.2 Native-token dependency as a first-order payments risk	24
2.8.3 Financial crime and sanctions perimeter: validator fee flows are not trivial	25
2.8.4 Singleness and “garden dynamics” under congestion.....	25
2.8.5 Cross-sector spillovers and coupling on shared public ledgers.....	25
2.8.6 Supervisory expectations if permissionless deployments are contemplated.....	26
2.8.7 Bank needs to be explicit about risk appetite, residual risk and accountability if permissionless deployment is contemplated.....	26
2.8.8 Future risks on these technologies?	26
2.9 Question 17-18: Comments or suggestions on Section 2.3 and 2.4	27
2.9.1 Redemption fees.....	27
2.9.2 Remuneration	28
2.10 Question 19: Wholesale, and non-sterling (Comments or suggestions on Section 2.5).....	29
2.10.1 Wholesale use of systemic stablecoin.....	29
2.10.2 Non-sterling stablecoin	32

1. Summary of response

Across our response, we treat the CP’s proposals as choices about the **architecture of sterling money and payments, not only issuer-level supervision. Our organising lens is *par production*: “£1=£1” is a performance promise, and it is only as credible as the system’s ability to deliver cash-now convertibility and settlement in normal times and under stress. This lens focuses attention on (i) where par is made in practice (issuer balance sheets, corridor design, settlement infrastructure, markets, and the central bank), (ii) where liquidity and operational constraints bind under scale and stress, how “cash-now” elasticity is sourced, and who bears the strain, and (iii) the resulting monetary footprint and incidence across banks and markets. The aim is not to eliminate stress margins, but to ensure they act as vital, transparent and governable discipline, rather than as hidden fragilities that fail discontinuously under stress.** This matters not only for singleness, but for monetary stability: the same design choices shape the pace and incidence of expansion/contraction in money-like claims, the terms on which settlement elasticity is provided, and—crucially for monetary sovereignty—the central bank’s ability to govern ‘cash-now’ settlement conditions in sterling in stress.

Viability “in the round” (not a policy goal, but a constraint on whether benefits materialise). We do not treat issuer profitability as a policy objective. But if the policy intent is to enable *safe sterling stablecoin innovation* (for productivity/competition/resilience gains), the regime should avoid being unnecessarily burdensome in ways that do not materially improve financial stability or singleness. Otherwise, the outcome can be a technically coherent regime that fails to attract participants—foreclosing potential benefits by design. The CP’s main explicit viability lever is the permitted backing mix; but viability is shaped cumulatively by a wider bundle of design choices, including holding limits, corridor frictions, non-remuneration, and the clarity/credibility of stress-liquidity arrangements. This response therefore repeatedly asks whether each burden is doing distinct stability work, or whether the same objectives could be met with lower-friction architecture.

Q1 (Backing assets): The consultation’s move away from a reserves-only posture toward permitting a material securities leg (up to 60% in short-term UK sovereign debt) is a regime change, not a minor “viability” tweak. It changes the *par-production mechanism, the binding stress margins, and the*

monetary footprint of a systemic stablecoin sector:

- **Par provision mechanism (issuer-level): reserves buffer plus monetisation capacity.** Under 100% BoE deposits, redemption is mechanically straightforward. Under hybrid backing, payments-grade par depends on stress-time monetisation capacity (speed, haircuts, market/settlement capacity, operating hours). The key evaluative point is that “safe assets” are not enough: a liquidity lens is essential, and the regime must be explicit about monetisation pathways and operating assumptions that make redemption reliable.¹
- **Two-way runs must be assumed.** Flight-to-safety inflows into a well-designed stablecoin are plausible (especially relative to uninsured deposits for larger holders). But outflows/redemption waves are also plausible, driven by rail risk (operational/governance incidents), spikes in off-chain settlement needs, or other persistent boundary needs. **Deleveraging provides a simple example:** if bank credit remains the locus of debt creation, loan repayment ultimately settles in bank-deposit/CB-money terms, meaning that even in a “successful” stablecoin world a deleveraging event will generate stablecoin→deposit redemption pressure.
- **System stress margins: reserves-only is asymmetric; hybrid makes stress more two-sided.** With reserves-only backing, flows out of coin are typically low-stress for banks (funding/liquidity return), while flows into coin are the binding risk (deposit/reserve drain, redistribution, potential LOLR reliance). Hybrid backing may moderate *steady-state* funding incidence for the banking system via partial recycling through asset purchases, but it does not remove the potential for acute short-run liquidity pressure on banks (conversion can run faster than recycling, and recycling is contingent and mediated by settlement cycles and market capacity). Hybrid backing also makes outflows from coin potentially stress-relevant for gilt/repo markets because correlated redemptions translate into time-critical cash demand against collateral (and potentially asset sales); at scale, this can resemble **MMF-style liquidity feedback loops** via repo demand/asset monetisation (“dash-for-cash” dynamics in safe assets).
- **Repo access is welcome and should be treated as core “par infrastructure” under hybrid.** Repo-based monetisation is the least destabilising route to “cash-now” relative to reliance on outright sales (especially where eligible assets are largely buy-

- to-hold). But repo is only as good as the terms and operational readiness—pre-positioning, eligible collateral, haircuts, hours, and the ability to execute at scale under congestion—recognising wrong-way risk (tightening precisely when cash-now demand spikes).
- **One-way repo protects issuer par, not system liquidity.** Allowing issuers to repo gilts to raise cash can improve redemption execution under hybrid backing, but it does not offset bank liquidity/funding pressures created by large inflows into stablecoin; and in stress it can add time-critical demand for repo cash even though redemption proceeds flow back into the banking system. The implication is not that issuers should become liquidity recyclers; rather, the regime should be explicit about the **cash-now stack** so payments-grade par is not left contingent on fragile market capacity. Allowing two-way repo would be a different regime choice—turning issuers into liquidity intermediaries participating in secured money markets, with MMF-/dealer-bank-like governance, limits, and backstop questions—beyond the narrow payments-issuer model the CP appears to envisage.
 - **Eligible-asset capacity and reverse repo as a possible conservative broadening. A narrow eligible-asset set can look conservative, but may create its own scalability and market-concentration risks. One possible conservative broadening would be to permit short-term reverse repo against high-quality gilt collateral, including longer-maturity gilts subject to conservative haircuts, rather than permitting outright holdings of longer-duration gilts. This could widen the effective collateral base while keeping the issuer’s direct asset short-term, secured and close to cash; reduce direct duration and price risk; lessen reliance on outright sales in slower asset markets; and attenuate some private-QE-like securities-absorption effects. It could also allow stablecoin inflows to supply cash into secured funding markets, rather than making the sector only a liquidity absorber. But it would be a material architecture choice: stablecoin issuers would become standing cash lenders in secured money markets, with implications for repo rates, collateral demand, haircuts, counterparty concentration, rollover risk, liquidity backstops and monetary-policy transmission.**
 - **Monetary footprint (full vs hybrid): cash/CBDC-like outside-money withdrawal discipline vs partial recycling with portfolio-channel effects (“private-QE-like”).** Under reserves-only backing, stablecoin issuance functions as a cash-like withdrawal into outside money: deposits are converted into a non-bank par claim backed one-for-one by central bank

deposits, imposing transparent funding/liquidity discipline on banks and making the monetary-policy and LOLR framework first-order. That incidence is not undesirable: bank funding and liquidity discipline is part of the quid-pro-quo for private credit creation; the policy aim is that the resulting pressures remain transparent and governable. Under hybrid backing, this **outside-money withdrawal channel** is moderated but still present (roughly in proportion to the reserves share), and an additional portfolio (“private-QE-like”) channel is introduced: by purchasing short-maturity UK government debt (gilts), often from non-banks, issuance can re-create deposits for sellers while maintaining a par transaction claim in circulation, increasing the stock of money-like liquidity held by the private sector even though the stablecoin issuer is not creating credit ex nihilo. At scale, this can loosen financial conditions via portfolio rebalancing and is therefore relevant to monetary transmission (and potentially price stability), making the reserves–securities split a macro design choice, not merely an issuer-viability tweak.

(See section 2.1).

Q2 (Step-up): We support step-up to reduce cliff edges and entry barriers, but the proposal creates a tension: allowing up to **95% gilts** early maximises reliance on **monetisation (repo/sales)** precisely when redemption capability is least proven. Step-up is credible only if “controlled operation” is made concrete—clear issuance/use-case constraints, observable readiness/stress-redemption triggers for scaling, a defined path to the steady-state mix, and a time limit to prevent step-up becoming a de facto permanent perimeter. Finally, step-up should be framed as **near-money scaling under constraints**, not permission to market “money-proper” usability with weaker liquidity: if it is used/advertised as general-purpose money, it must meet money-grade par and operational standards, or remain visibly constrained. (See section 2.2).

Q4 (Trust / ring-fencing): We support holding backing assets and required reserves on statutory trust (insolvency-remoteness is foundational for a payments-grade par claim), but emphasise that “trust” is not a single design: outcomes depend on custody/control, reconciliation, waterfall clarity, and funded operational continuity in wind-down. We also flag a semantic risk: over-emphasising “beneficial ownership” can pull public understanding toward an MMF/NAV “owning gilts” mental model; the regime should frame trust as a mechanism to deliver a uniform par monetary claim, with clear yield treatment and tight governance of wind-down reserves. (See section 2.3).

¹ For a non-technical discussion of why liquidity and cash-now capacity (not only solvency) are the binding constraints for payments-grade stablecoins, see Sansom (2025a, 2025b).

Q8–10 (Holding limits): Holding limits are best understood as a prudential tool with architectural consequences. They can be defensible as a transitional guardrail—while authorities and market participants learn about adoption dynamics and stress behaviour—but they are not “free safety”: they reshape incentives, operating models, and where burdens sit.

- **Stock→flow + mode switch:** Limits are framed as a *stock* cap to contain disintermediation/run risk, but if they bind in routine payment circuits they become a *flow* architecture choice: either brittle failures (rejections/queues and edge-case fragility) or reliance on sweeping/waterfalls that turn stablecoin into a corridor-dependent rail, with high-frequency coin↔deposit churn. In that rail mode, boundary performance (speed, fees, hours, capacity, conduct) becomes critical infrastructure and a potential gatekeeper—so payments-grade corridor standards matter as much as the headline cap.
- **Perimeter/enforceability (“who is counted”).** Meaningful limits require identity aggregation across wallets/providers (and potentially across issuers). Without it, limits are porous and incentivise fragmentation/routing. This forces an explicit perimeter choice: a regulated identity/wallet layer capable of aggregation, or acceptance that limits can only bite reliably at regulated gateways (issuance/redemption), with leakage where transfers occur outside that perimeter - and a further choice about where enforcement sits (on-chain vs intermediated), with different governance and operational-resilience burdens.
- **Alternatives/complements to holding limits (first lines of defence):** stock caps are a blunt tool for run management (which is about *speed* and time-critical liquidity). A more robust defence-in-depth is: (i) reduce the store-of-value pull into stablecoins (non-remuneration and conduct/marketing constraints); (ii) make coin↔deposit interoperability payments-grade by construction (service-level standards on price/speed/reliability, with transparent rule-based queuing/windows under load); and (iii) specify stress-time liquidity arrangements and monetisation pathways (repo readiness, operating-hours assumptions, and—under hybrid backing—credible backstops). Holding limits can then play a transitional containment role, rather than carrying the full burden of stability. (See section 2.5).

Q11 (Direct access): We support requiring systemic issuers to access interoperability-relevant payment systems *directly* rather than via a sponsoring participant. Sponsorship can economise on intraday liquidity through netting/timing transformation,

but it does so by re-introducing tiering and sponsor-dependent limits/terms at the most sensitive boundary; direct access helps keep coin↔deposit convertibility rule-based and reduces concentration and gatekeeper risk at the corridor. Hybrid backing strengthens incentives to economise on reserve buffers and reduce reliance on stress-time monetisation—making sponsored netting/timing-liquidity attractive, but also making corridor dependence more salient. Direct access should be framed as “no sponsor dependence” (not “no service providers”): issuers may use agents for connectivity, but settlement balances and redemption performance should not hinge on discretionary sponsor capacity. (See section 2.6.)

Q14–16 (Public permissionless ledgers): We broadly agree with the Bank’s top-level risk buckets (accountability, finality, operational resilience), but stress that on permissionless rails the binding risks are often **execution and governance dependencies external to sterling**: (i) reliance on a volatile native token and fee market to obtain inclusion (“gas” as an execution constraint), (ii) shared-blockspace coupling to unrelated activity, and (iii) protocol-level counterparties (validators) that sit outside UK accountability—each of which can impair payments-grade convertibility and practical singleness under stress even when backing assets are pristine.

- **Collapsed bearer model versus two-layer model.** The risks are strongest where the stablecoin is treated as a native bearer instrument whose legal, monetary and operational life is collapsed into the public ledger itself. A regulated two-layer model, in which an issuer or Token Operator maintains the legally recognised holder record while using the public ledger as an evidentiary and execution layer, may mitigate asset-safety and record-reconstruction concerns. But it does not remove finality risk; it relocates it into the recognition rulebook. The Bank should require clarity on recognition thresholds, pre-recognition use, downstream reliance, correction powers, migration, liability and residual loss allocation. For payments-grade settlement and collateral use, unrecognised balances should not be spendable, pledgeable or reusable as final value unless a regulated intermediary explicitly bears the resulting settlement or credit risk.
- **Native-token (“gas”) dependency is a first-order payments risk, not a footnote.** Even if a stablecoin is fully backed by central bank money, core functions (on-chain transfer, mint/burn/redemption, emergency controls) typically **depend on a volatile non-sterling execution asset** to secure inclusion. Under congestion, fee markets can reprice

endogenously and inclusion becomes uncertain, creating a new instability channel for payments-grade par: *delay* → *discount* → *confidence/run dynamics*, even when backing assets are pristine. If sterling transfers require a non-sterling native token to execute, then blockspace—and the native-token liquidity that buys it—becomes the scarce execution resource that gates on-chain finality and can delay practical convertibility at par, weakening the idea that central bank money is unconditionally “top” in operational terms. If central bank money is to remain operationally top, fees/capacity allocation must be sterling-native and governable (or insulated via reserved capacity/deterministic inclusion), rather than mediated by a volatile external token and fee market—because a core attribute of monetary sovereignty is the central bank’s ability to govern “cash-now” settlement conditions in its own currency in stress.

- **Sanctions/AML perimeter and externalities are intrinsic.** Protocol-level fee payments route value to validators/block producers who may be pseudonymous and outside UK accountability; KYC’d stablecoin users do not remove counterparty exposure at the protocol layer. At scale, a systemic sterling instrument can also subsidise the wider ecosystem via fee flows and native-token demand, raising questions about acceptable residual risk and required design constraints.
- **Singleness risks: scale-dependent lock-in, with boundary operations bottlenecking first under stress.** Even in normal conditions, if boundary crossing (coin↔deposit and cross-issuer conversion) is slower/costlier/less reliable than on-rail transfers, users rationally avoid the boundary where they can (“garden dynamics”), weakening practical singleness—and the incentive strengthens as on-rail acceptance reaches critical mass (network effects). **Under congestion, boundary operations are typically the most complex and “must-clear” steps (mint/burn, compliance checks, payout orchestration), so the corridor can become the first bottleneck—even if peer-to-peer transfers continue for those willing to transact on-chain—making sterling interoperability state-contingent and potentially encouraging issuer/platform clustering (issuers gravitate to the dominant “garden” to minimise boundary friction).**
- **Shared blockspace creates cross-sector coupling (unlike dedicated payment rails).** On a public chain, payments-critical functions compete for inclusion with unrelated activity (DeFi liquidations, bridge incidents, memecoin bursts). That means “execute-now capacity” for sterling payments is exposed to shocks originating outside the payments perimeter. This contrasts with current arrangements where

payments rails and market infrastructures are segregated and capacity/governance are purpose-built; on a shared chain, equivalent robustness would require explicit segregation mechanisms (reserved capacity / deterministic inclusion policies / credible fallbacks), otherwise coupling is a design choice with consequences for whether “payments-grade” service levels are realistically deliverable.

- **Implications for policy and supervision.** If permissionless deployments are contemplated, the Bank should state risk appetite and accountability up front, and make dependencies measurable and governable: stress tests that include gas price spikes/fee-market dysfunction/validator disruption, explicit continuity targets (time-to-finality/time-to-redeem), clear incident playbooks (including fork/reorg policy), and an explicit stance on how regulated entities must manage (or may tolerate) protocol-level sanctions/AML exposure. (See section 2.8).

Q17 (Fees / redemption): We support the Bank’s stability framing on redemption fees, but the CP’s drafting (Bank of England, 2025b, Section 2.3.3) risks conflating **payment-service fees** with **frictions on exercising a par monetary claim**. Fees are integral to payment chains, but charging for (or delaying) conversion of £1 coin into £1 deposit is not simply a service fee: if material or state-contingent it functions as an exchange-rate-like toll within sterling, creating garden dynamics, encouraging secondary-market exit and away-from-par pricing, and can amplify confidence/run dynamics (and “garden” lock-in). For a payments-grade instrument, a cleaner principle is therefore a presumption of **free (or de minimis) redemption via at least one standard route**, with simple ex-ante terms, **no stress-time widening**, and **no loss pass-through**—while leaving ample scope to charge for genuine processing/acceptance and value-added services. Regulation should also anchor to the **user-experienced boundary outcome** (all-in price, speed and reliability across the full corridor), not the label attached to a charge. (See section 2.9.1).

Q18 (Remuneration): We welcome continuity with the 2023 baseline and continue to support non-remuneration as an anchor for payments-grade money, keeping systemic stablecoins oriented toward transactional/operational balances rather than yield-seeking savings (while recognising that relative attractiveness can still vary with the interest-rate environment, especially for corporate cash managers). We agree remuneration would be regime-defining: if stablecoins (directly or equivalently) pay interest, they become a rate-sensitive portfolio asset competing with deposits/MMFs, enlarging the run surface and

raising monetary-transmission questions via the spread between stablecoin and deposit returns. (See section 2.9.2).

Q19 (Wholesale): We welcome the Bank’s anchor and low risk appetite for any material shift away from wholesale settlement in central bank money, and strongly support public-rail routes (RT2 modernisation, synchronisation/omnibus models) as the most direct way to deliver new functionality while preserving singleness. The key question is not “private money vs CB money”, but where CB money sits in the stack—as the **final settlement anchor**, the **interoperability/corridor anchor**, and the **stress-liquidity anchor**. This matters because widespread wholesale stablecoin use could **displace banks’ stable funding** while the system still needs **time-critical ‘cash-now’ elasticity** that today is largely supplied by banks and, ultimately, the central bank—creating a potential tension unless the elasticity layer is explicit and governed. A practical binding constraint is therefore the settlement-liquidity stack (netting/LSMs, queuing, intraday credit, collateral monetisation): wholesale settlement is not “RTGS everywhere”; it relies on netting/LSMs and intraday credit to economise on scarce settlement balances. Stablecoin rails are naturally prefunded/RTGS-like unless they explicitly rebuild that stack—there is no third option (either accept prefunding, or reintroduce credit/monetisation mechanisms with clear governance and backstops). Proposals to “rebuild better” via high-frequency netting and an explicit credit-for-settlement layer may be coherent, but once intraday liquidity is supplied and priced at scale it becomes monetary-policy-adjacent (corridor/floor-ceiling design, incentives, and gridlock externalities become first-order unless the design keeps intraday liquidity strictly separable from overnight funding). We emphasise two-way run risk (deposit/MMF→token vs token→deposits/CB money) and that the backing model determines where stress binds (bank funding/liquidity distribution and corridor capacity under reserves-heavy backing; gilt/repo monetisation capacity and market-liquidity spillovers under hybrid backing). Holding limits do not scale cleanly to wholesale. The core design choice is therefore an explicit endpoint: a bounded settlement scheme versus a money-like wholesale cash asset (which requires explicit corridor standards, monetisation pathways, and credible stress-liquidity arrangements able to absorb both run directions). (See section 2.10.1).

Q20–22 (Non-sterling stablecoins): Non-sterling stablecoin risk should be assessed by functional dependence and infrastructure role, not treated as a single “foreign coin” problem. For ordinary UK domestic activity, currency-substitution risk appears comparatively low because sterling is embedded in a

mutually reinforcing system of price stability, public obligations, sterling wages and debts, safe bank money, accounting and pricing conventions, and liquid sterling payment and settlement infrastructure. The most important defence against currency substitution is therefore a credible, liquid and resilient sterling monetary system, not a weakening of money-grade safeguards in order to promote sterling stablecoins. The more plausible risk is infrastructure-dependence: in cross-border payments, tokenised securities, collateral movement, treasury management, wholesale settlement, margining and platform commerce, non-sterling stablecoins may become important because liquidity, composability, market depth and network effects are organised around a foreign-currency claim. London may legitimately host valuable multi-currency activity without that implying domestic dollarisation; the policy concern is whether sterling-relevant markets, UK-facing users, or time-critical payment and collateral functions become dependent on foreign-currency liquidity pools, reserve assets, governance arrangements or conversion corridors. Deference to home authorities may be appropriate for issuer-level prudential safeguards, but UK-facing corridor accountability remains necessary for operational resilience, conduct, FX transparency, data visibility, interoperability, incident response and treatment of UK users in stress. Further two-country balance-sheet analysis would be valuable because denomination, issuance, use, reserve investment and FX/corridor intermediation may sit in different jurisdictions. (See section 2.10.2).

2. CP Questions responses

2.1. Question 1: Backing assets

Q1: *Do you have views on our proposal to allow systemic stablecoin issuers to hold up to 60% of backing assets in short-term sterling-denominated UK sovereign debt securities alongside unremunerated deposits at the Bank, as an appropriate balance between business model viability and mitigation of financial stability risks?*

2.1.1. Why allowing securities investment is not just a business-model tweak

The proposed 40/60 calibration should not be viewed as a narrow “business model viability versus safety” trade-off at the level of an individual issuer. The choice between full central bank deposit backing vs a hybrid of central bank deposits plus short-dated gilts is, in effect, a **system architecture choice. It changes (i) the technology of par—from direct liquidity (reserves) to partial reliance on rapid monetisation of securities; (ii) the monetary footprint—from a more cash/CBDC-like deposit-substitution channel toward a regime in which reinvestment can recycle reserves back into the banking system with portfolio effects ; and (iii) the locus of stress—shifting vulnerabilities from bank-side funding and liquidity distribution toward the functioning and capacity of the gilt/repo complex (and the boundary between payments and markets). For that reason, the appropriate question is not only whether a given issuer can meet redemptions in steady state, but whether the backing rule, monetisation pathway and operating assumptions together deliver payments-grade par convertibility under congestion and stress**, without importing market microstructure fragilities into core payments (Sansom, 2025b).

2.1.2. Benchmark: what 100% BoE deposits buys (and what it shifts onto banks)

The 100% backing requirement in unremunerated deposits at the Bank of England provides the cleanest “par technology” from the perspective of the stablecoin itself. Redemption is funded by direct central bank liquidity rather than by selling or financing securities, so the issuer’s ability to meet predictable and correlated outflows is not exposed to market liquidity conditions, haircuts, dealer balance-sheet constraints, or settlement-cycle frictions. In that sense, reserves-only backing minimises the risk that stress in collateral markets feeds back into the payment instrument’s ability to convert at par.

However, this does not eliminate system-wide incidence; it relocates it. A systemic stablecoin cannot expand

“ex nihilo” in the way bank deposits can through credit creation. Growth in the stablecoin float occurs by converting bank deposits into a non-bank sterling claim, with an associated transfer of reserves to support backing at the central bank. Deposit creation within the banking system can be relatively “self-financing” when credit expansion is broad-based and payment flows are largely offsetting across banks so net interbank reserve settlement needs remain small relative to gross flows (Kumhof & Salgado-moreno, 2024). By contrast, growth of reserve backed stablecoin is mechanically a one-way leakage from deposits into a non-bank par claim, **so it necessarily generates** net settlement and refinancing pressure for the banking sector. The banking system must then either replace that funding through wholesale markets and/or central bank facilities, or adjust balance sheets (including, at the margin, the supply of credit).

This refinancing margin need not be destabilising—even in a scarce-reserve environment—but it is the channel through which the monetary policy operating framework and lender-of-last-resort architecture become first-order. The central bank, as monopoly supplier of reserves, can accommodate on terms consistent with policy objectives; in that sense, a larger non-bank par claim can increase the sensitivity of bank funding conditions to the policy stance, potentially strengthening monetary policy traction. In an ample-reserves regime, the effect may primarily be to absorb excess reserves (supporting balance-sheet “re-normalisation”), but it can still tighten conditions at the margin through distributional reserve scarcity and refinancing frictions of the type emphasised in the “money channel” (Kumhof & Salgado-moreno, 2024) (i.e., what matters is not just aggregate reserves but where they sit and how easily they can be reallocated).

The broader point is that reserves-backed systemic stablecoin would have a cash-like footprint: it can be a scalable non-bank sterling instrument that displaces deposits. That footprint can support payments utility, and monetary policy discipline, but comes with run/flight-to-safety dynamics risk. The stability of the overall arrangement therefore leans on familiar foundations across the existing system — payments-grade convertibility of commercial bank money into central bank money, robust money-market redistribution mechanisms, and credible lender-of-last-resort capacity—rather than on the new issuer’s asset-liability management.

2.1.3. What the hybrid 40/60 backing changes

The proposed 40/60 hybrid changes the regime in economically material ways. Most importantly, it

changes the **technology of par**: once a significant share of backing assets is held in short-dated gilts, the issuer’s ability to deliver rapid redemption at par depends not only on the stock of central bank deposits it holds, but on its ability to **monetise backing assets at speed and at scale**. In normal conditions this can be achieved via repo or outright sale, but under stress it becomes contingent on repo market capacity, haircut dynamics, dealer balance-sheet constraints, settlement calendars and operating hours. This introduces a distinct vulnerability: dysfunction in the gilt/repo complex—or simply congestion and operational timing constraints—can translate into delayed or more costly redemption, undermining confidence in par precisely when the instrument is meant to be most reliable.

A key implication is that **outright sales are not a robust “elasticity” mechanism** for a payments-grade instrument (promising par on demand and fast payments) facing correlated redemption pressure (whether via portfolio adjustments or directional payments flows). Importantly, correlated outflows need not be confidence-driven: so long as bank credit remains the locus of debt creation and repayment, deleveraging episodes require settlement in bank-deposit/central-bank-money terms, generating stablecoin→deposit conversion pressure that—under hybrid backing—will be met in practice via rapid gilt monetisation.

Even if sales can work in steady state, they can be too slow at the point of stress (and can be destabilising at scale through price impact), creating a two-way spillover problem: redemption waves can become a source of stress in gilt/repo markets, while gilt/repo stress can impair redemption performance and import market fragilities into payments. In practice, **stress-time resilience hinges on repo-based monetisation (or equivalent liquidity facilities)**—so credible repo access and an operational playbook for using it are not an optional refinement but essential for credibility of “par” and singleness under a hybrid regime.

Allowing reinvestment in securities also changes the **monetary and banking-system incidence** of stablecoin growth. Under a pure reserves model, issuance is a near-mechanical **sink** for bank deposits and reserve liquidity and source of re-financing/funding risk for banks: £100 of deposits converted into coin becomes £100 of coin backed by £100 of central bank deposits, reducing commercial bank deposits by £100. This is a “cash like” monetary footprint (see 2.1.2). Under a 40/60 hybrid, only the 40% reserves share is parked at the central bank, while 60% is funded by purchasing gilts (including from non-bank sellers) and paying sellers via the banking system (returning reserves and deposits to banking system).

This is closer in spirit to “private QE” than to deposit substitution: it can increase the stock of liquid par claims held by the private sector (**coins plus re-created deposits**), even though it does not create money via bank-style credit extension - in a simple stylised case, £100 of deposits converted into coin becomes £100 of coin **plus** £60 of commercial bank deposits held by the gilt sellers (all retail means of payment/purchasing power), with the net banking sector deposit drain reduced to £40 – thus the hybrid model can moderate steady-state pressure on aggregate bank funding. In addition, the gilt leg can create a **portfolio-rebalancing channel**: by purchasing short-dated gilts—especially from non-bank holders—the issuer increases demand for eligible sovereign paper, puts downward pressure on yields at the front end, and can induce broader reallocation into other assets.

However, whilst steady-state funding pressure may be moderated, it does not remove the potential for acute short-run liquidity risk: recycling via asset purchases is not a fast response to redemption spikes, and if issuers are unable to deploy liquidity via reverse repo, they can still behave as a short-run liquidity sink within the system in stress. This raises a distinct design question: should repo be treated only as a stress-time monetisation tool, or also as an ordinary backing asset through short-term reverse repo? A regime limited to outright holdings of very short-dated gilts may be conservative in credit and duration terms, but may also be non-scalable if the eligible asset bucket is small relative to plausible stablecoin issuance. It may also concentrate demand in a narrow part of the gilt market. A possible conservative broadening would be to permit short-term reverse repo as an eligible backing asset, collateralised by high-quality gilts, potentially including longer-maturity gilts subject to conservative haircuts. In that structure, the issuer’s direct asset is a short-term secured cash claim rather than an outright longer-duration gilt exposure.

However, this should be recognised as a regime choice, not a technical relaxation. Treating reverse repo as an ordinary backing asset would make systemic stablecoin issuers standing cash lenders in secured money markets. The sector would not merely monetise gilts in stress; it would allocate stablecoin proceeds into repo markets in normal times. This may improve scalability and reduce concentration in a scarce short-gilt bucket, but it also changes the monetary footprint of the regime by adding a secured money-market participation channel alongside the portfolio channel from outright securities purchases.

The policy distinction is therefore between: (i) repo as a stress monetisation tool, where issuers borrow cash against assets they already hold; (ii) reverse repo as an

ordinary eligible backing asset, where issuers lend cash into secured money markets against collateral; and (iii) outright investment in broader or longer-duration securities. These are not equivalent. The first supports par under stress; the second makes stablecoin issuers part of the secured money-market ecosystem; the third introduces more direct market-risk and portfolio-balance channels.

For these reasons, the 40/60 model moves systemic stablecoin away from a simple “reserves wrapper” with cash/retail CBDC like footprint, and toward a government-MMF-like liquidity transformation structure (but a stricter par promise and payments-grade par convertibility implies continuous, operationally critical liquidity on demand) with private-QE like footprint. Even with high-quality, short-duration assets, it concentrates dependence in the front-end gilt and repo markets and raises the bar for the operating assumptions needed to make par credible. The policy implication is that “asset safety” alone is not sufficient: if the Bank permits a material securities leg, it should specify the expected monetisation pathway and operating assumptions that make payments-grade par reliable under congestion and stress—particularly in a world where a systemic payments instrument may be expected to function beyond traditional market hours. This is precisely why liquidity arrangements and backstops become first-order, not a footnote. At the same time, the securities component introduces a **portfolio-rebalancing channel** closer in spirit to QE: purchases of short-dated gilts can compress yields in the eligible bucket and, via portfolio rebalancing, propagate more broadly (see 2.1.6).

2.1.4. Achieving “£1 = £1 everywhere”: interoperability at the coin↔deposit boundary

If systemic stablecoins are to operate as a general-purpose sterling payment instrument—rather than a closed-loop scheme—then singleness requires low-friction interoperability between stablecoin wallets and bank accounts; the relevant question becomes how the coin↔deposit corridor is designed and operated.

Singleness of money is experienced at the boundary, not in the asset schedule. Even if backing assets are high quality, users only experience “£1 = £1 everywhere” if coin↔deposit conversion used in routine payment journeys is near-par, predictable, and operationally reliable. It is therefore important to distinguish between (i) *redemption* (“I want my money back”) and (ii) *payments-driven conversion* (“I need to pay someone on the other rail”), which can generate high-frequency corridor flows mechanically through merchant sweeping, payroll, and treasury processes. If payments-driven conversion is slow, costly, or uncertain, it is experienced as an exchange-rate-like friction within sterling, creating

incentives to remain on-rail (“garden dynamics”) or to rely on intermediaries to smooth the boundary.

In practice there are two corridor modalities. One is high-frequency settlement in central bank money, in which corridor flows are settled frequently in reserves; this preserves singleness but raises operational capacity and intraday liquidity intensity, and—under a hybrid backing model—makes “payments-grade par” depend on how reliably securities can be monetised into reserves at speed and scale. The other is warehousing/float at the boundary, in which intermediaries net and smooth corridor flows and redeem on a schedule; this reduces continuous settlement load but recreates intermediation at the most system-sensitive point (timing mismatch, operational dependence, and potentially sponsor-like reliance). Either way, the policy implication is that backing rules and corridor design must be specified as a package: expected time-to-convert at par in routine journeys (not only a legal redemption right), out-of-hours operating assumptions, stress-time operational standards, and a fee/term framework that avoids the boundary becoming a persistent “toll booth” inside sterling.

2.1.5. Why repo access is needed and reverse repo should be considered.

If the regime permits a material securities component, **repo-based monetisation is essential**: it is the mechanism that allows the issuer to generate cash quickly without relying on potentially destabilising asset sales and therefore sits at the heart of payments-grade par under stress. This is especially so given the proposal to limit stablecoin issuers to investment in very short-run government debt that is largely a buy to hold market where you would not want to rely on deep secondary-market liquidity as a stress redemption strategy.

It is also important to recognise the **directional footprint** of the proposed permission set. If issuers can repo securities to obtain cash but cannot deploy cash via reverse repo, the stablecoin sector is structurally more like a **liquidity absorber** than a recycler: inflows park reserves while holding securities, and outflows monetise securities for cash to meet redemptions. In stress, that can add to secured funding demand without offsetting supply, potentially tightening marginal repo conditions for banks/dealers. This does not mechanically “break par” (the central bank can accommodate), but it affects where adjustment pressure and non-linearities appear—another reason to specify operating assumptions and backstop expectations explicitly.

The case for considering reverse repo is best understood as a conservative-widening and liquidity-architecture argument, not as a general relaxation of asset quality.

Properly constrained, short-term reverse repo against high-quality gilt collateral could broaden the usable collateral base without requiring all non-reserve backing to sit in a narrow short-gilt bucket. It could also reduce direct duration and price risk for issuers: the issuer’s asset would be an overnight or short-term secured cash claim, rather than an outright holding of longer-duration gilts. This matters because simply widening outright eligible assets could push issuers into buying and selling assets in markets that may be slower, thinner or more volatile than is compatible with payments-grade “cash-now” convertibility under stress. Reverse repo may therefore provide a way to improve scalability without relying on outright sales in less liquid asset markets or exposing issuers directly to broader duration risk.

Reverse repo may also reduce some asset-market impact relative to a broader outright asset-backed model. Outright purchases and sales by large stablecoin issuers could amplify price pressure in the eligible asset bucket, especially if issuance and redemption flows are large or correlated. Reverse repo changes that mechanism: the issuer supplies cash against collateral rather than simply purchasing and removing securities from market portfolios, or selling them under stress. In that sense, reverse repo may attenuate the private-QE-like channel associated with direct gilt purchases (Sansom et al., 2026). Stablecoin issuance would still create a money-like liability, but the asset-side adjustment would shift from outright securities absorption toward secured funding conditions.

This does not make the regime neutral. It changes the channel. A one-way repo model lets issuers raise cash against gilts when they need redemption liquidity; reverse repo would also let issuers supply cash into secured funding markets in normal times, so the stablecoin sector would be less purely a liquidity sink. That may be more governable than broad outright asset eligibility, because repo markets have established control variables — collateral eligibility, haircuts, margining, counterparty limits, clearing or tri-party arrangements, maturity limits and concentration limits. But it would also make those repo-market controls part of the stablecoin regime itself. The Bank would therefore need to specify the terms under which systemic stablecoin issuers could become cash lenders in secured money markets, including counterparty eligibility, collateral treatment, encumbrance, settlement timing, stress testing, rollover assumptions and interaction with any central bank liquidity backstop.

This boundary is not merely theoretical. The US payment-stablecoin framework includes certain

overnight reverse repurchase agreements within the eligible reserve-asset perimeter, with the issuer acting as purchaser of securities, collateralised by Treasury securities, and subject to overcollateralisation and specified tri-party, centrally cleared or bilateral counterparty conditions (Regulation of Payment Stablecoins, 2026). That contrasts with the UK proposal as read here, where repo appears primarily as a route to monetise backing assets rather than as an ordinary reverse-repo backing asset. This divergence is useful because it highlights the design choice: should systemic stablecoin issuers be permitted only to raise liquidity against securities, or also to supply liquidity into secured money markets?

2.1.6. Liquidity backstop essential for payments grade par under hybrid backing

We note the CP confirms that the Bank is considering a potential liquidity backstop for systemic stablecoin issuers. This is the right question to confront explicitly: once a material share of backing assets is held in short-dated gilts, payments-grade par becomes contingent on stress-time monetisation capacity (repo access, haircuts, dealer balance-sheet, settlement cycles and operating hours). In normal conditions that capacity may be ample; in correlated redemptions or gilt/repo dysfunction it may not be. If the policy objective is rapid, predictable convertibility at par under stress, a credible backstop (or equivalent committed liquidity arrangement) is therefore not an optional enhancement but an essential part of the “par technology” itself. Conversely, if no backstop is envisaged, the regime should be explicit that the instrument is a bounded scheme whose usability in stress is managed through designed frictions (e.g. redemption windows, liquidity fees/gates) rather than guaranteed on-demand convertibility.

A useful benchmark is government money market funds: a short-maturing sovereign portfolio can support same-day redemption in normal conditions, but recent experience shows that even high-quality, short-duration structures can face liquidity externalities in stress and may require official liquidity support to avoid fire-sale dynamics (Bank of England, 2021; Hauser, 2021; Kashyap, 2020). Recent episodes—most clearly the March 2020 “dash for cash” (BCBS-CPMI-IOSCO, 2021; FSB, 2020; Kashyap, 2020) and the UK gilt/LDI stress of September–October 2022 (Alexander et al., 2023; Bank of England, 2022; Gabor, 2023; Giuzio et al., 2024) — illustrate that when many actors simultaneously need “cash now”, liquidity constraints and forced selling can dominate even where underlying assets are high quality, prompting policy moves to introduce official backstops (FSB, 2025). The introduction of the Fed’s MMF Liquidity Facility in March 2020, and subsequent

regulatory reforms (US Federal Reserve, 2020; US SEC, 2023), underline that “safe assets” do not remove the need for a credible stress-liquidity mechanism when redemptions become correlated. The Bank’s own CNRF (Bank of England, 2025a) similarly illustrates the direction of travel: in severe gilt-market dysfunction, non-bank liquidity needs can become systemically relevant and may warrant central bank lending against gilts to reduce forced sales. This is further underlined by ongoing concern about fragilities in government bond backed repo markets (FSB, 2026). However, the CNRF is contingent and eligibility-limited; it is not designed as a standing mechanism for payments-grade convertibility. The policy implication is therefore that— if hybrid backing is permitted at systemic scale—the regime should specify up front how stress liquidity will be accessed (through whom, against what collateral, on what terms, and with what operating hours), and how tail risk is allocated between the issuer, intermediaries, and the central bank.

The MMF analogy should be used carefully. MMFs are near-money: investors may redeem because they need settlement money. A payment stablecoin with MMF-like backing could differ if it is itself accepted for payment and settlement, because holders would not necessarily need to redeem merely to obtain settlement media (WBS Gillmore Centre Financial Technology, 2024). That may reduce one source of MMF fragility, but it does not remove the liquidity-transformation problem. It changes its location: MMF-style liquidity transformation would now sit inside the payment-money layer. Repo-market resilience, collateral eligibility, haircut dynamics and backstop design would therefore become part of the infrastructure supporting sterling par itself.

2.1.7. Macro and monetary footprint (why the sign is ambiguous and driven by 60/40).

The macro effects of a hybrid backing regime are not signed a priori. A reserve-backed component has a cash/CBDC-like footprint: when bank deposits are converted into stablecoins backed by central-bank deposits, the banking sector loses deposits and reserves, and any replacement funding must come through banks’ own balance-sheet adjustment or central-bank liquidity provision. By contrast, a securities-backed component can recycle cash back into the banking system. If the issuer buys securities from banks, banks exchange securities for cash, although the original deposit funding has still been displaced. If the issuer buys securities from non-banks, the seller receives a bank deposit and the seller’s bank receives reserves. This can moderate the aggregate bank-funding and reserve-drain effect, but it is also the case in which stablecoin issuance most closely resembles private QE: the stablecoin remains outstanding, the non-bank seller

now holds a bank deposit, and securities have been absorbed from market portfolios.

The reserves–securities split therefore does not merely calibrate issuer liquidity. It shapes the balance between a more cash/CBDC-like deposit-substitution channel and a more private-QE-like securities-absorption channel. This is why the regulatory cap matters in practice. Issuers are likely to choose the highest-yielding permissible mix, so absent explicit expectations on liquidity buffers, rebalancing cadence and stress-time mobilisation, the effective macro footprint may sit at or near the regulatory boundary. The net effect will depend on scale, holder mix and the operating framework, but the locus of adjustment follows from the backing calibration. This architecture-first point is developed more fully in our related stablecoin monetary-implications paper (Sansom et al., 2026). Reverse repo would add a third channel. If issuers are permitted to lend cash against collateral as an ordinary backing asset, the non-reserve component need not operate through outright securities absorption. Stablecoin inflows would instead create a supply of cash into secured funding markets. Lending to banks or dealers could recycle reserves back to the banking system and allow central-bank refinancing to be repaid, but it would not recreate the displaced deposit base: banks would instead be funded through secured wholesale borrowing from the stablecoin issuer. Lending to non-banks could still expand broad private transaction money, because the stablecoin remains outstanding while the non-bank borrower receives a bank deposit. But the mechanism would be short-term secured lending, not outright asset purchase.

Reverse repo may therefore be less private-QE-like in the narrow portfolio-balance sense, and potentially more governable. The issuer holds a short-term secured claim rather than an outright securities portfolio, while the borrower generally retains economic exposure to the collateral. The adjustment margin shifts from securities scarcity, yields and portfolio rebalancing toward repo rates, eligible collateral, haircuts, margining, maturity, counterparty limits, clearing or tri-party arrangements and rollover conditions. This does not make the regime neutral. It changes the monetary footprint from outright asset absorption to secured money-market intermediation.

The UK monetary-framework developments reinforce why this should be treated as a monetary-architecture question rather than passive investment plumbing. As the Bank transitions toward a repo-led, demand-driven Sterling Monetary Framework, reserves, collateral and secured funding markets become central to the operating environment in which systemic stablecoins

would exist. The Bank’s work on gilt repo market resilience and the CNRF point in the same direction: collateralised cash lending is part of financial-stability and monetary architecture.²

The policy implication is that backing rules should be assessed across three distinct margins: reserve warehousing, outright securities absorption and secured lending against collateral. A hybrid reserve/securities regime may reduce one set of pressures while creating another. Reverse repo may attenuate the private-QE channel and make the adjustment margin more governable, but it would make systemic stablecoin issuers part of the secured funding architecture. Whilst the proposal to restrict stablecoin issuers from reverse repo is a defensible conservative boundary, the consultation’s treatment of repo therefore also matters since that is a monetary-design choice, not simply an eligible-asset question.

2.1.8. Non-remuneration anchors payments use; remuneration opens portfolio channel

We support non-remuneration as a sensible anchor for a payments-grade instrument: it helps keep systemic stablecoins oriented toward transactional and operational liquidity needs rather than becoming a wholesale savings substitute, reinforcing a more cash-like monetary footprint. We recognise the point (made by (Bindseil & McLellan, 2025; Ulrich Bindseil, 2025) that prohibiting remuneration does not automatically enhance stability and may create unintended cyclicalities by making relative attractiveness vary with the level and path of interest rates. However, this likely overstates the opportunity-cost channel for payments-motivated holdings: transactional balances and operational buffers are typically already minimised (sized to “not fail payments”), so marginal changes in the policy rate mainly affect issuer economics and the size of the yield spread rather than inducing large swings in required demand for stablecoin balances.

By contrast, if systemic stablecoins become remunerated they could cease to be primarily a payments float and could become a rate-sensitive portfolio asset. They would then compete with bank deposits and MMFs for corporate cash management, making demand more fast-moving and interest-elastic. In that regime, a return differential between stablecoins (via backing yields and pass-through, with a differential policy rate) and bank deposits could become a meaningful policy lever that shifts balances between deposits and coins; the flip side could be a

materially larger disintermediation and run-risk surface. Given a non-remuneration requirement, permitting a material securities component means the issuer earns the yield on eligible short-dated gilts while coinholders earn zero, creating a visible spread (“private seigniorage”) with political-economy salience. Banks earn a spread in part because deposit creation is tied to balance-sheet risk and refinancing discipline; a reserves-only stablecoin does not create money ex nihilo and under 100% reserve backing warehouses central bank money in a private wrapper, so any spread in that model looks less like a policy-created rent for operating the wrapper (when payment services value added would be more naturally recovered via transparent, competitive fees). Allowing a securities component reintroduces liquidity transformation and market-monetisation risk; backing income then has a clearer economic role in funding buffers and stress monetisation capacity, but also creates classic incentives for risk-creep unless governance and backstops are explicit.

For these reasons, there is a strong case for transparency on how backing income is used and for guardrails that align incentives with resilience (liquidity buffers, operational readiness, and wind-down funding) rather than encouraging minimum liquidity or lobbying for riskier eligible assets (expanded under Q13/Q4).

2.2. Question 2: step-up regime

Q2: Do you have comments on the step-up regime as a way of supporting innovation while mitigating financial stability risks?

We agree with the motivation for a step-up regime. It can reduce cliff-edge effects, lower barriers to entry, and avoid unintentionally favouring incumbents who can reach scale quickly because of network effects. In that sense, step-up can support innovation by allowing operational capability and business models to be demonstrated before the full “systemic-grade” burden applies.

However, the proposed design creates a central operational tension. Permitting up to 95% short-dated gilts at the earliest stage may be pragmatic for business viability, but it also implies the least immediately liquid backing and therefore the greatest reliance on monetisation (repo/sales) (see 2.1.3) precisely when an issuer’s liquidity-management capability and redemption operations may be least tested. This places substantial weight on “controlled operation” and

mobilisation limits (Bank of England, 2025b, p. 17) to contain risk—so the credibility of step-up depends on those controls being clear, enforceable, and observable. For that reason, we encourage the Bank to specify what “controlled operation” means in practice, at least directionally. In particular: (i) initial mobilisation/issuance limits and permitted use cases; (ii) the metrics and triggers for increasing limits (e.g. growth, concentration, redemption performance under stress tests, operational readiness); (iii) the expected pathway and criteria for stepping down from 95% gilts toward the steady-state mix; and (iv) the maximum time an issuer can remain in step-up absent meeting readiness criteria. A step-up regime without explicit transition discipline risks becoming a de facto long-run regime.

A helpful way to interpret the perimeter is to distinguish “money-proper” from “money-like” or “near-money”, rather than “systemic” vs “non-systemic”. The CP’s recognition that an issuer could be “systemic at launch” reflects that the relevant trigger is not just balance-sheet size but qualitative moneyiness—whether the instrument is likely to function as general-purpose sterling money from the outset. This matters because users and merchants will not reliably discriminate between coins operating under different perimeters if both are marketed as “£1 = £1” and widely spendable. Step-up therefore works best when it is understood as “scaling under constraints” for near-money experimentation, rather than as permission to offer money-like usability while operating with materially weaker liquidity.

The implication is that step-up should not “step up” the user-facing money standard. Even if a coin is small, if it is used and perceived as money, a failure to redeem at par can generate confidence spillovers at the category level (i.e. mistrust of “stablecoins as money”). If backing liquidity is materially weaker during step-up, the regime should require compensating constraints so users and counterparties do not treat it as fully money-grade—for example tighter limits and/or narrower use-case scope, stronger disclosure/branding that it is operating under a constrained runway, and readiness tests that are prerequisites for expanding usability.

Finally, step-up should be linked explicitly to the par-production mechanism. During the early stage, par is produced less by a cash buffer and more by monetisation capacity; the Bank should therefore explain how step-up interacts with permitted repo operations, the stress assumptions around market capacity and haircuts, and the intended role (if any) of any Bank liquidity backstop (emergency-only versus operationally usable). A constructive way to operationalise this is a two-dimensional step-up: scaling not only with size but with

demonstrated operational readiness (tested redemption operations, pre-positioned collateral, proven ability to raise reserves under stress). In short, we support step-up as an innovation enabler—but only if it is coupled with clear containment, explicit transition discipline, and an unambiguous commitment that once a coin becomes meaningfully money-like in usage, it must meet money-like standards of par reliability, redemption capacity and operational resilience.

2.3. Question 3-6: Risk-based capital and reserve requirements

Q3: Do you agree with our approach to mitigating risks to the issuer and coinholders via risk-based capital and reserve requirements? If not, what approach would you see as more appropriate for systemic stablecoin issuers?

Not addressed in this response.

Q4: Do you agree with our proposal that the reserves of liquid assets to mitigate the financial risk of backing assets and cost of insolvency/wind down should be held on trust ring-fenced from the general estate of the issuer? If not, do you have alternative proposals to mitigate risks to coinholders in the event of issuer failure/insolvency and in the absence of a specific set of arrangements to deal with failure?

2.3.1. “Trust” model good, but semantic risk can pull understanding toward an investment-share mental model (“owners not creditors”)

We agree in principle with the proposal to hold backing assets and required reserves on a statutory trust, ring-fenced from the issuer’s general estate. If the policy objective is a payments-grade par instrument rather than a fund-like claim, insolvency-remoteness is a necessary foundation. Properly designed, a trust structure strengthens coinholder protection in issuer insolvency and supports confidence that redemption will be honoured at £1 = £1.

That said, “trust” is not a single design. The key question is not trust versus no trust, but whether the legal and operational mechanics deliver *full and timely* payout in messy failure states. Outcomes will depend on the detailed structure (including custody and control), the waterfall, operational continuity, and the ability to validate records, reconcile entitlements, and mobilise assets for distribution under stress.

We welcome the CP’s clarity that the proposed model gives coinholders a beneficial (proprietary) interest in safeguarded assets held on trust, supplemented by a residual unsecured claim against the issuer to the extent of any shortfall. However, **there is a public-**

² Against that UK background, the Federal Reserve’s ON RRP facility is a useful comparator (Federal Reserve Bank of New York, 2026). It illustrates how large-scale institutional cash-lending arrangements can affect money-market rate architecture and private repo-market activity (Hempel et al., 2023).

facing semantic risk: over-emphasising “trust” can pull understanding toward an investment-share mental model (“owners not creditors” / “owning gilts”) (one case in point (Rhodes, 2025)). That interpretation is unhelpful for a payments instrument (a **statutory trust is naturally “asset-side language”** (beneficial interest in a pool), while a payments stablecoin is meant to behave like a **liability-side object** (a uniform, par, on-demand monetary claim)). The regime should therefore frame the trust consistently as an insolvency-remote mechanism for *par convertibility*, not as a claim on specific assets in an MMF/NAV sense. Disclosures should avoid implying beneficial ownership of particular backing assets or any NAV/yield exposure; the core message should remain that holders have a par monetary claim and the ring-fenced pool is the mechanism to support redemption and orderly wind-down.

We also support the proposal for an insolvency/wind-down reserve held on trust, but only if the governance is tight. If the trust structure is intended to substitute for a full resolution regime, it must include credible funding for the “plumbing” of wind-down—records validation, reconciliation, custody access, asset liquidation or repo, distributions, and critical operational continuity. For credibility and to preserve coinholder protection in practice, eligible costs should be narrowly defined, spend should be capped and controlled, independent assurance/audit should be embedded, and the waterfall should be explicit (including when and how the wind-down reserve can be used).

Finally, custody and control standards matter as much as the legal wrapper. For non-BoE deposits and securities holdings, the regime should require qualified third-party custody, strict segregation, robust daily reconciliation and discrepancy resolution processes, operational resilience requirements, and independent audit/assurance. Outsourcing custody should not dilute accountability: the issuer should remain responsible as trustee for ensuring assets are controlled, identifiable, and readily mobilisable in both normal operations and failure states.

Given that eligible gilts will generate income, the regime should also clarify the “who gets the yield?” rule. It should be clear whether and how backing income can be used to fund operations and required reserves, what must be retained to support resilience and wind-down readiness, and what disclosures apply—so incentives are aligned with safety outcomes and coinholders are not exposed to opaque value transfer.

If the Bank’s long-run objective is bank-like “equivalent outcomes” for systemic issuers, an explicit special

administration or resolution-style regime (with continuity powers) may ultimately be a necessary complement, so that the trust is not the only pillar carrying failure management.

In practical terms, the proposal would be strengthened by clarity on: (i) minimum operational standards for segregation, reconciliation and audit; (ii) expected payout speed for coinholders in key failure scenarios (service-level targets); and (iii) the permitted use and governance of the insolvency/wind-down reserve and the waterfall ordering.

Q5: Do you have views on our proposal for calibrating capital for general business risk?

Not addressed in this response.

Q6: Do you have views on calibrating the reserve requirements for insolvency/wind down?

Not addressed in this response.

2.4. Question 7: Other comments or suggestions on Section 2.2

Q7: Do you have any other comments or suggestions on the proposals on the major policy revisions set out in Section 2.2?

Although fees are discussed in Section 2.3 of the CP, we treat redemption/conversion frictions as system-relevant because they condition interoperability and par in stress; we therefore comment in detail under Q17 (see Section 2.9).

2.5. Question 8-10: Holding limits and the risks they address

Holding limits are often presented as a simple safeguard against disintermediation and run risk. In practice, they are also an architectural lever. Enforceable limits require choices about identity, wallet perimeter and where controls sit (on-chain vs regulated intermediaries). And—crucially—tight limits for businesses change how the instrument can be used: because many economically meaningful payment journeys have a business endpoint (merchant receipts, payroll funding, B2B settlement), binding business limits tend to force rapid sweeping back into bank deposits. That shifts the system footprint from a “stock” effect (persistent balances outside deposits) toward a “flow” effect (high-frequency stablecoin↔deposit conversion), concentrating operational and liquidity demands at the corridor where singleness is

experienced (availability, speed, fees, operating hours, and intraday liquidity/float). Q8 and Q9 therefore need to be read together: the practicalities of implementing limits and the usability of stablecoins under those limits are two sides of the same design choice—money-like instrument versus corridor-dependent rail.

Q8: What are the operational challenges to implementing holding limits or other tools we are exploring? How might those challenges be addressed, including for individual and business limits?

2.5.1. Operational challenges (and why they’re not just technical issues)

Holding limits are often presented as a straightforward safeguard against disintermediation and run dynamics. In practice, implementing limits at scale forces a set of operational and perimeter choices that shape what a systemic stablecoin *is*: a money-like instrument that can circulate on its own balance sheet, or a corridor-dependent “rail” that must continually convert back into bank deposits. The challenges are therefore not only technical; they go to enforceability, usability, where risk and operational load sit in the system, and whether stablecoin can (i) provide a genuine outside option that disciplines payments pricing and access terms and (ii) shift the monetary incidence by creating a scalable outside money like constraint on bank funding/credit conditions at the margin, or whether tight limits instead re-intermediate it into a corridor product priced and governed at the bank boundary.

2.5.2. Identity and aggregation: the hard problem is not the limit, it is “who is counted”

Limits are easy to apply **per wallet address**. They are **materially harder to apply per person or per business across multiple wallets, providers, and (if limits are per coin) multiple issuers**. Without robust identity linking and aggregation, limits become porous: **users can split balances across wallets or issuers, weakening the policy objective and incentivising fragmentation and complex routing**. Addressing this requires an explicit perimeter choice—either a **regulated-wallet/identity layer that can link multiple wallets to a beneficial owner, or acceptance that enforceable limits can only reliably bite at regulated boundaries (e.g. issuance/redemption points), with leakage risk where transfers can occur outside that perimeter**.

2.5.3. Where enforcement sits: on-chain strength versus governance and resilience burdens

A second challenge is **where** the limit is enforced and what that implies for governance and operational resilience. **On-chain enforcement** (e.g. contract-level caps) can be robust in principle, but typically implies

upgradeability and administrative controls to handle exceptions, recovery, legal orders, migrations, and incident response. Those features introduce governance risks and increase operational-resilience obligations, because the enforcement mechanism becomes part of the critical payments infrastructure. **Off-chain enforcement** (issuer or regulated-wallet controls) is often more practical, but it depends on consistent implementation across providers and—crucially—on reliable identity mapping. It also raises circumvention questions if unhosted wallets and on-chain transfers remain broadly permitted, because users may be able to route around regulated providers.

2.5.4. Failure modes and user experience: limits without “waterfall” are hard constraints

In the absence of an explicit, mandated “waterfall” (for example, automatic sweep into a linked bank account), holding limits will operate as **hard constraints**: transactions are rejected, queued, or otherwise fail once a limit is hit. That makes limits not just a prudential guardrail but a binding usability feature. A credible regime therefore needs deterministic, ex ante rules for predictable edge cases: what happens if a user hits the limit mid-payment; how refunds and chargebacks are handled; how limits interact with batching, merchant settlement timing, and payroll cycles; and what happens during predictable high-flow periods. These are not minor implementation details: ambiguity over rejection versus queuing, partial acceptance, or “grace bands” can itself become a stress amplifier if users and counterparties cannot predict outcomes at the point of need.

2.5.5. How holding limits can turn “money” into “rails”

The most consequential issue is how business limits interact with real-economy payment patterns. In modern economies, **businesses are the endpoints for a large share of payment flows**—merchant receipts, payroll funding, supplier payments, and routine B2B treasury operations. If businesses can hold sufficient working balances, stablecoin can function (at least in part) as a money-like settlement balance: firms can accept receipts, manage liquidity within operational cycles, and choose when to convert in batch. If, however, business limits bind below working-balance needs, stablecoin becomes usable for commerce only if receipts can be swept out rapidly and reliably into deposits. In that regime, the system’s design burden shifts: the key question becomes less “how safe are the backing assets?” and more “how robust is interoperability under load?”—including availability, speed, fees, and operating hours, and who provides intraday timing liquidity/float when conversion cannot be perfectly synchronised. In effect, tight business limits are an architectural pivot: they push

stablecoin away from an asset-like monetary layer that can circulate, and toward a corridor-dependent rail whose viability rests on payments-grade conversion infrastructure.

Q9: *What are your views on the usability of stablecoins in the presence of holding limits, both for individuals and businesses? What use cases do you envisage would require exemptions from the proposed limits? What uses would not be possible given the proposed limits?*

2.5.6. Usability depends on whether stablecoin is “money” or a corridor-dependent “rail”

Holding limits do not just cap balances; they shape the *mode of operation*. If limits bind routinely—especially for firms—stablecoin can remain usable only if the stablecoin↔deposit corridor delivers *payments-grade conversion*: fast, predictable, low-friction, and operationally reliable under peak loads and (where relevant) outside market hours. If corridor performance is weak or uncertain, users will respond rationally by either (i) holding precautionary buffers to avoid failed transactions or (ii) avoiding boundary crossing altogether (“garden dynamics”), which undermines singleness and reduces the practical usefulness of stablecoins as a general-purpose sterling instrument. In other words, tight limits are coherent only alongside a credible interoperability stack. Without that, limits risk producing a brittle user experience (failed payments at the margin) and, systemically, converting what is presented as a money-like instrument into a conversion-heavy rail whose viability and pricing are dominated by the bank-linked boundary.

2.5.7. Individuals: what remains feasible, and what really drives usability

The CP’s proposed retail holding limits (e.g., £20k per coin) are framed as compatible with high usability for most individuals, supported by indicative metrics (such as coverage relative to salary and typical current-account balances). We agree that, in a steady-state sense, many day-to-day retail payment use cases could remain feasible within that range.

However, the binding determinant of usability is not only the headline cap. It is whether users can rely on conversion at the boundary such that they *do not need to hoard balances* to insure against frictions. Even modest frictions—uncertain settlement windows, unclear fees, or operational downtime—create an incentive to hold larger precautionary balances *within* the stablecoin rail to avoid getting stuck mid-journey (especially if payments must land in bank accounts). That behavioural response can erode the intended effect of limits while also increasing the salience of “garden dynamics” (users stay on-rail to avoid boundary tolls and failures).

The practical implication is that, for individuals, holding limits and corridor service levels are complements: if limits are to be binding without becoming disruptive, the corridor must be sufficiently smooth that users can run low balances and still trust that routine conversion for payments will succeed quickly and predictably.

2.5.8. Businesses: the decisive constraint, and why exemptions are not edge cases

For businesses, holding limits are not simply “larger retail limits.” Firms sit at the endpoints of high-volume payment circuits: merchant receipts, platform payouts, payroll funding, supplier payments, and routine B2B treasury processes. If households use stablecoins, businesses must be able to accept them; if businesses use stablecoins, they must be able to hold and deploy them operationally.

Two further considerations sharpen the point. First, cash-management incentives—especially under non-remuneration—mean firms will typically want to sweep excess non-interest-bearing balances into interest-bearing instruments anyway. So, a degree of “payments float + sweeping” is a natural equilibrium even without binding limits. Second, what matters operationally is whether limits sit *above* or *below* these natural working-balance needs for a firm. If limits are above working balances, firms can accept stablecoin receipts, manage intraday inflows/outflows, and choose when to sweep. Stablecoin can function (at least partially) as a settlement balance.

If, instead, limits are set below working operational needs, sweeping stops being discretionary and becomes *forced and high-frequency*. At that point stablecoin ceases to be a balance-sheet settlement asset for the business sector and becomes usable only as a transient rail that must continually convert back into deposits. The banking system’s burden shifts from a “stock” effect (persistent stablecoin balances displacing deposits) toward a “flow” effect: constant churn through the corridor and intraday liquidity implications from this, with usability now dominated by corridor capacity, fees, operating hours, and the provision of intraday timing liquidity/float when conversions cannot be perfectly synchronised.

This is why exemptions are not an edge case but a core design choice. To keep routine commerce viable under binding business limits, the regime will likely need principled, rule-based exemptions (or differentiated operational accounts) for entities that must warehouse operational liquidity to process flows—e.g., PSPs, acquirers, large merchants, treasury hubs, or scheme-like intermediaries. But exemptions are not costless: they can (i) concentrate operational and

systemic importance in a small number of boundary intermediaries, and (ii) become a backdoor route to wholesale stablecoin holdings unless tightly defined, monitored, and constrained by purpose (operational liquidity, not investment/cash management) and by governance/controls.

The policy question is therefore not “should there be exemptions?” but “which actors must be exempt for the payments system to function under the chosen limits—and how do we prevent exemptions from simply relocating large holdings into a smaller set of warehousing intermediaries?”

2.5.9. What uses would be structurally discouraged or impossible under binding limits

If holding limits bind for businesses (or bind tightly in normal times rather than only in stress), several use cases become difficult or structurally discouraged:

- **Large merchant ecosystems holding receipts through intraday or settlement cycles** without continuous sweeping. Acceptance becomes viable only if conversion is near-automatic and highly reliable.
- **Corporate treasury use of stablecoin as a settlement balance**, except where exemptions apply. Stablecoin can no longer serve as an operational store of settlement liquidity if balances must be kept below working needs.
- **“In-coin” closed payment loops at scale** (where receipts fund onward payments within the stablecoin rail). Binding limits force the system back into deposits quickly, so stablecoin behaves less like circulating money and more like a transfer wrapper around deposit rails, but with implications for bank settlement modality and intraday liquidity needs.
- **Any model that relies on stablecoin balances as a buffer against operational timing mismatches** unless the corridor provides equivalent buffering via speed and reliability of conversion, or intermediated float.

These are not marginal losses: they define whether stablecoin is a distinct settlement asset or primarily a corridor-dependent messaging/execution layer.

2.5.10. How holding limits shape competition and gatekeeping

Holding limits also shape competitive dynamics. In a rail-like regime (forced sweeping), the corridor becomes the key price bottleneck: conversion fees/terms/capacity/hours matter more than on-chain transfer costs in determining the end-user experience and the effective price of using stablecoins. That creates a risk that banks and/or a small set of regulated intermediaries become de facto gatekeepers—able to

mute competitive pressure through pricing, throttling, or operational terms at the boundary.

This has two system-level implications. First, it can blunt stablecoin’s ability to act as a genuine outside option that disciplines payments pricing and access. Second, it can relocate the banking-system stress margin: tighter limits may reduce the *structural* stock displacement effect on bank funding, but they can increase *intraday* liquidity and operational capacity demands at the corridor—precisely where throttling and pricing power can be exercised.

For that reason, a holding-limit regime that is intended to preserve usability and competition should be paired with clear expectations for corridor performance and conduct (service levels, operating hours assumptions, transparency of fees/terms, and stress-time operating rules), so that limits do not inadvertently turn “competition in payments” into “competition in corridor pricing” governed by a small number of boundary providers.

Q10: *Other than holding limits, what do respondents consider are the tools best suited to mitigating the risks we have identified?*

2.5.11. Tools other than holding limits (and how to position them)

Holding limits can be a useful supplementary guardrail, but they are not the only—or necessarily the most robust—way to mitigate the risks identified in the CP. In our view, it is helpful to organise the alternative toolkit by function: tools that **reduce the “store-of-value pull”** into stablecoins (so the instrument remains payments-oriented), and tools that **make the conversion corridor safe by construction** (so interoperability does not become a hidden fragility or a gatekeeper bottleneck). Where any “throttling” is contemplated, it is better treated explicitly as a **flow management** mechanism at the conversion boundary rather than implicitly via blunt stock constraints.

2.5.12. Clarifying what holding limits are trying to achieve in this CP

We read the CP’s motivation for holding limits primarily as (i) reducing the risk of **large and rapid outflows from deposits into a new non-bank sterling claim**—especially under stress—and (ii) limiting the potential for such shifts to interact with bank funding conditions and credit provision while the market and authorities learn about adoption dynamics. That objective is conceptually distinct from day-to-day payments usability. It matters because different objectives point to different tools: if the concern is mainly **payments competition and innovation**, the priority is frictionless and well-governed

interoperability; if the concern is mainly **stress-time flight-to-safety**, the priority is managing the *pull* and the *pace* of conversion under stress.

2.5.13. Why stock caps are a blunt instrument for run dynamics

A holding limit caps an end-state stock. Run dynamics, by contrast, are typically about **speed and coordination**: rapid conversion flows can be destabilising even if the eventual stock is bounded. In practice, a cap can still permit a very large conversion wave up to the cap, while simultaneously creating brittle edge cases (rejections/queuing) and incentives to route around limits (multi-wallet and multi-issuer splitting). For these reasons, stock caps can be a weak substitute for a clearer and more targeted approach to either (i) reducing the store-of-value attractiveness of stablecoins or (ii) managing conversion flows transparently at the boundary.

2.5.14. Reduce the store-of-value pull (preferred first line of defence)

If the policy aim is to ensure systemic stablecoins remain **payments-grade instruments** rather than wholesale savings substitutes, the cleanest tools are those that reduce the incentive to hold stablecoins for yield or investment purposes. **Non-remuneration is therefore an important anchor (as discussed above – see 2.1.8): it pushes stablecoin demand toward transactional and operational balances, rather than rate-sensitive portfolio allocation. This can be reinforced by perimeter and conduct tools—such as clear restrictions on marketing stablecoins as savings products, and expectations that fee structures and product design do not replicate remuneration indirectly through rebates or loyalty-style yield equivalents. The principle is simple: the more stablecoins behave like money for payments, the less the regime must rely on brittle balance caps to contain store-of-value demand.**

2.5.15. Make the corridor payments-grade by construction (avoid boundary fragility)

A key alternative to relying on holding limits is to ensure that interoperability between stablecoin wallets and bank accounts is **payments-grade**: fast, predictable, operationally resilient, and transparent on fees and terms. Concretely, this points to tools such as (i) explicit expectations on time-to-convert at par in routine journeys (service-level style commitments), (ii) operational availability assumptions (including how out-of-hours demand is handled), and (iii) transparent fee principles so that boundary crossing does not become a persistent “toll booth” inside sterling. This approach mitigates two risks at once: it reduces incentives for users to hoard balances to avoid boundary frictions,

and it reduces the scope for gatekeeper dynamics where corridor providers can mute competitive pressure through pricing or throttling.

2.5.16. Flow tools at the boundary (if throttling is contemplated at all)

If the regime judges that some form of stress-time throttling is required, it is more coherent to use **explicit flow tools** than to assume a static holding cap will act as a withdrawal throttle. Examples include deterministic queuing rules, transparent settlement windows, and clear, ex ante commitments on how conversion requests are processed under congestion. The key requirement is that any such tools must be **rule-based and pre-specified**: discretionary or opaque throttling is confidence-sensitive and can itself become a run trigger. The aim should be to make outcomes predictable under load, rather than creating uncertainty about whether and when “£1 converts to £1”.

2.5.17. Perimeter and structural tools (enforceability, interoperability, and transition discipline)

Several tools are best understood as perimeter choices that determine what can be enforced and where. If policy intent is to enforce limits per beneficial owner, that implies an identity and aggregation architecture (and associated standards for wallet providers). If the aim is to avoid fragmentation and routing arbitrage, interoperability standards and common operational requirements across issuers and wallets become more important than per-issuer caps. Finally, the CP’s stepped or transitional structure makes “migration discipline” a tool in its own right: clear expectations for how coins and issuers move from an FCA-supervised perimeter to the Bank’s systemic perimeter can reduce cliff effects and user confusion while preserving the integrity of the “money-like” standard once broad usability is reached.

2.5.18. Note on other limit types considered in the CP

We note the CP also considers alternative limit designs (e.g. aggregate caps and transaction limits) and acknowledges important drawbacks, including the risk that some designs could induce deviations from par or create undesirable frictions in routine payments. We agree with that caution. Tools that are opaque to users or that make convertibility contingent in unpredictable ways risk undermining singleness and confidence. Where constraints are used, it is preferable that they are transparent, rule-based, and paired with a clear interoperability stack—so that the regime mitigates risks without quietly degrading the money-like character of the instrument.

2.6. Question 11: Direct access

Q11: Do you have views on our proposal that systemic stablecoins should access payments systems that support interoperability across different forms of money directly rather than through a sponsoring participant?

*We support the direction of travel: for a systemic stablecoin, direct access to payments systems that support interoperability across forms of money is best understood as a **monetary-architecture safeguard**, not merely a choice about operational efficiency. The policy objective is that coin↔bank-money conversion and boundary-crossing within sterling feels **near-par, fast, and reliable**—not like an exchange-rate-like friction that encourages “walled gardens”. Designing interoperability as a *system property* is more robust than making it a *sponsor privilege*.*

The core trade-off is liquidity efficiency versus tiering and dependence. Sponsorship can economise on intraday liquidity because a sponsor can pool flows across clients and smooth timing via queuing/netting and, in practice, sometimes balance-sheet intermediation (even if only intraday). But those liquidity “savings” are not neutral: they arise because the issuer’s par promise becomes partly mediated through a **private counterparty at the most systemically sensitive boundary**. That reintroduces a form of tiering—coinholders’ ability to move back into bank money depends on the sponsor’s limits, operational continuity, and willingness to act—precisely where the regime is trying to eliminate discretionary frictions.

This matters directly for the Bank’s own stability logic. In stress, a sponsoring model can turn sponsor constraints (risk limits, liquidity hoarding, outages, de-risking, or termination) into **delay/friction at the redemption boundary**, which can then propagate through the familiar channel: delay → secondary-market discount → confidence loss/run dynamics. Direct access reduces the probability that par convertibility depends on discretionary private balance-sheet capacity and a concentrated set of gatekeepers at the corridor.

That said, “direct access” should be framed as **no dependence**, not “no service providers.” Systemic issuers may still use technical agents for connectivity and operational services, but the regime should ensure that:

- settlement balances remain **the issuer’s own** (not a sponsor liability or pooled float),
- payment submission and redemption performance are **not subject to discretionary sponsor gating** (pricing/terms/availability),
- any liquidity support is **explicit, secured, and portable** (not implicit reliance on a sponsor’s

balance sheet).

We also encourage the Bank to clarify the implementation expectation. In particular:

- Is direct access a **hard requirement at authorisation**, or a condition to be met within a defined transition window as the issuer scales?
- What does “direct access” mean operationally (e.g., RTGS/CHAPS settlement account and ISO 20022 connectivity; any role for other UK payment rails), and what are the implied expectations around operating hours and resilience?
- How does the direct-access policy interact with the broader liquidity model (including any “one-way” monetisation arrangements), so that the route to obtaining reserves for redemption does not simply reappear as a new sponsor bottleneck elsewhere in the chain?

Overall, we see direct access for systemic issuers as a key step toward keeping singleness and interoperability robust under scale and stress: it reduces concentration and gatekeeper risk at the corridor, and makes convertibility a rule-based system outcome rather than an outcome contingent on a small set of private intermediaries.

2.7. Question 12-13: Backing asset safeguarding and proposed trust arrangements

Q12: Do you agree with our proposed approach to safeguard backing assets? If not, what alternative measures do you propose?

Not addressed in this response.

Q13: Do you have views on the proposed legal structure of the trust arrangements for backing assets and reserves to deliver the desired outcomes set out in this consultation paper? This includes feedback on the overall structure of the trust arrangements and whether these should be structured as a single trust covering both backing assets and reserve requirements or as two or more separate trusts.

Evaluate single trust vs. multiple against clear objective: “£1 = £1” in practice, including in wind-down

The “single trust vs multiple trusts” question should be evaluated against the desired outcomes. The key test is not the legal form in isolation, but whether the arrangements deliver: (i) clear beneficiary rights and priority (who gets paid, in what order, and for what purpose), (ii) speed and certainty of payout in messy failure states (including the ability to mobilise assets under stress), and (iii) operational simplicity and auditability (so the structure is legible to run, supervise, and enforce, with low litigation/ambiguity risk). In

other words, the trust architecture should be judged by whether it reliably produces “£1 = £1” in practice, including in wind-down.

This links directly to the point we made under Q4 (see 2.3.1): a statutory trust can strengthen coinholder protection and insolvency-remoteness, but “trust” is not a single design. Outcomes depend on custody/control, reconciliation, the waterfall, and—critically—whether the structure explicitly funds the plumbing of payout (records validation, reconciliation, asset monetisation, distributions, and continuity of critical ops). A single trust that looks strongest on paper can still fail the policy objective in practice if it does not solve, ex ante, who funds and executes the mechanics when the general estate is impaired.

Avoiding the semantic trap: trust as *par-promise safeguarding*, not “gilt ownership”

As flagged in Q4, there is a public-facing risk that “assets held on trust” gets misread as “users are owners not creditors” or “coinholders own gilts” (e.g. Rhodes, 2025). That interpretation is economically misaligned with a payments instrument. The trust should be framed consistently as an insolvency-remote mechanism to secure performance of a par monetary promise (supported by the CP’s own clarification that coinholders have a beneficial interest in safeguarded assets plus a residual unsecured claim to the extent of any shortfall), not as an MMF-style unit claim with NAV/ yield exposure. This framing choice matters because it shapes expectations and behaviour in stress.

A workable default: one statutory trust with hard internal segregation and an explicit waterfall?

A single statutory trust that contains clearly segregated sub-funds and a pre-agreed waterfall, for example:

- Backing pool (assets held to support par redemption),
- Financial-risk reserve (to mitigate asset/operational risks consistent with the CP),
- Wind-down / administration reserve (to fund orderly payout and critical continuity).

To be credible, the trust deed (and associated supervisory expectations) should specify:

- what wind-down / administration costs are eligible,
- caps and controls on spend (including independent validation/audit),
- the order of priority between (a) funding orderly payout/continuity and (b) distributions to coinholders,
- “no discretion in stress” principles where appropriate (to avoid state-contingent dilution of the par promise).

This structure reduces legal fragmentation while still

delivering economic segregation and makes supervision and audit simpler: one legal wrapper, but internal ring-fences that can be tested and monitored.

When multiple trusts may be justified

A two-trust structure can be justified if the Bank wants maximum clarity that wind-down funding is insulated and immediately available even when failure is messy, for example:

- **Trust A:** backing assets (and potentially the financial-risk reserve) established primarily for coinholders’ protection;
- **Trust B:** a dedicated wind-down reserve, with tightly limited permitted uses, potentially structured so the insolvency practitioner/administrator can access it only for the narrow purpose of funding orderly payout/continuity.

The trade-off is that multiple trusts can reduce ambiguity about purpose and access in distress, but at the cost of greater operational/legal overhead and potentially more edge-case complexity (especially around triggers, governance, and cross-trust coordination).

Make yield treatment and conflicts explicit

With gilts in the eligible set, the backing pool generates income. Given the non-remuneration anchor for coinholders (and the CP’s acceptance that issuers may earn returns on backing assets), the trust structure should make the handling of income explicit so it strengthens—rather than weakens—the par promise. In particular, we recommend the regime clearly establishes that backing income should first support:

- reserve building and operational resilience,
- credible wind-down readiness (including tested payout capability),
- and only thereafter any distributable surplus (if any), subject to conditions that redemption performance and stress-time monetisation capacity are not impaired.

This reduces the incentive for “risk creep” and helps align governance with payments-grade outcomes.

Operational enforceability matters as much as legal architecture

Whatever structure is chosen, the Bank should treat minimum operational standards as part of the trust design, not an implementation afterthought: qualified third-party custody, strict segregation, frequent reconciliation with defined discrepancy handling, independent assurance/audit, and tested playbooks for rapid payout and asset mobilisation (including under stress). The objective should be that the trust arrangement is not only legally insolvency-remote, but operationally mobilisation-ready.

Keep the system legible to users

Finally, whichever structure is adopted, disclosures and drafting should keep the user-facing promise clear and consistent: holders have a par monetary claim; the trust is the safeguarding mechanism for that claim; it is not an invitation to treat the coin as “gilt ownership” or a NAV-exposed investment product. This clarity is itself stabilising—because it reduces scope for expectation-driven stress dynamics.

2.8. Question 14-16: Public permissionless ledgers

Q14: Do you agree with the Bank’s view that the prominent risks around public permissionless ledgers are accountability, settlement finality, and operational resilience, including cybersecurity?

We broadly agree with the Bank’s framing.

Accountability, settlement finality, and operational resilience/cybersecurity are the right top-level buckets for assessing public permissionless ledgers in a payments context, and they interact: weak or diffuse accountability makes it harder to manage incidents that directly affect finality or operational continuity when time is critical.

That said, it is worth avoiding an implication that these risks are unique to permissionless systems. Many payment rails face resilience risks (telecoms/ internet dependence, outages, cyber attacks). The distinguishing feature of permissionless ledgers is not that operational disruption is possible, but that **governance/accountability and execution economics are structurally different**: control is not concentrated in a regulated operator, prioritisation is mediated by fee markets and validators, and incident response can be constrained by the ledger’s own governance and client ecosystem.

2.8.1. Collapsed bearer model versus operator-recognised two-layer model

A further architectural distinction is important. The Bank’s concerns about permissionless ledgers are strongest where a stablecoin is treated as a native bearer instrument whose legal, monetary and operational life is collapsed into the distributed ledger itself. In that model, failure or ambiguity of the public ledger can become failure or ambiguity of the money. If the chain forks, reorganises, stalls, censors transactions or produces an ambiguous technical state, there may be no separate regulated record that can determine the legally recognised ownership position.

However, this is not the only possible architecture. In a two-layer model, a public ledger might be used as an evidentiary and execution layer by a regulated issuer

or Token Operator responsible for maintaining the legally recognised holder record (Ram, 2026b, 2026a). In such a model, a ledger outage, fork, reorganisation or congestion event need not by itself compromise the reserve assets, the legal claim or the ability to reconstruct the recognised record. The public ledger supplies evidence of attempted or recognised events; the legally determinative record is maintained under a regulated rulebook.

This changes the location of risk rather than eliminating it. If the Operator’s recognised record is legally determinative, the public ledger is not the ultimate ownership ledger in the legal sense. It is an evidentiary, reconciliation and possible programmable-execution layer beneath an operator-governed legal record. The key regulatory question then becomes how raw technical ledger events become legally recognised events, and when that recognition is final for downstream reliance.

This point is directly relevant to any reform of settlement-finality legislation for DLT-based systems. We welcome the Bank’s recognition that settlement-finality legislation may need to be updated to account for DLT-based systems. However, legislative recognition of DLT finality should not be framed simply as a bridge from “technical finality” to “legal finality”. The harder question is what legally protected system is being constituted. Traditional settlement-finality regimes protect transfer orders, netting and collateral within identifiable systems governed by rulebooks that specify entry, irrevocability and settlement finality. A public permissionless ledger does not by itself supply the same legal perimeter.

If a regulated Token Operator model is used, finality legislation should therefore attach to the operator-recognised system state, recognised event types and rulebook-defined moments of entry, irrevocability and finality, rather than to raw public-chain inclusion as such. Otherwise, the regime risks giving legal effect either too early, by protecting provisional technical state, or too vaguely, by failing to specify which state downstream parties may safely rely on. This is consistent with the broader legal challenge identified in settlement-finality reform: DLT-based systems require the law to identify a legally meaningful point of finality in environments where technical finality and legal finality may not naturally coincide (Löber, 2026).

The rulebook therefore becomes central. It should specify which ledgers and event types are recognised; the confirmation or recognition threshold; how forks, reorganisations, outages and data-availability failures are treated; the consequences of acting before operator

recognition; and the Operator’s powers to freeze, correct, replace, migrate or otherwise recover the recognised record. For smart-contract use cases, the contract, oracle or interface should be able to identify whether it is relying on raw chain inclusion or on an operator-recognised final state. Otherwise, apparent composability may rest on provisional technical evidence rather than legally final value.

This point is not merely conceptual. Once stablecoins are used in wider transaction chains, the stablecoin issuer can control only the stablecoin leg. It will not necessarily have authority to unwind the other legs of transactions that have already settled, including asset transfers, collateral movements, liquidations, hedges or downstream smart-contract executions. If an unrecognised stablecoin balance is used to trigger delivery of another asset, release of collateral, or execution of a hedge, and the Operator later refuses recognition or corrects the stablecoin record, one party may have delivered final value against a stablecoin balance that was never legally final.

For payments-grade settlement and collateral use, unrecognised balances should therefore not be spendable, pledgeable or reusable as final value. Raw chain inclusion may provide evidence of an attempted transfer, but operator recognition should be the point at which the recipient obtains a legally usable balance. If a wallet, platform or intermediary wants to provide instant usability before operator recognition, that should be treated explicitly as credit, guarantee or settlement-risk intermediation by a regulated party, with clear capital, liquidity, conduct and loss-allocation arrangements. It should not be allowed to emerge silently through smart-contract reliance on provisional balances.

A two-layer architecture may therefore turn open permissionless composability into governed composability. That may be appropriate for systemic or wholesale use cases, but it requires explicit rules on recognition, correction, liability, dispute resolution, migration and residual loss allocation. It also requires the public verification surface to remain aligned with the operator-recognised state. Where the Operator freezes, corrects, replaces or migrates balances, those interventions should be visible, where feasible, on-chain or through a verifiable state-publication mechanism, so that the public record does not become a misleading parallel history.

In these terms, the relevant policy question is not only whether a public ledger can achieve technical finality. It is whether the stablecoin arrangement creates sufficient legal and operational closure for other obligations to rely on it. A transfer may be valuable, technically included and even reconstructible from public data, but

still not be suitable for settlement or collateral use if its legal recognition remains provisional. This distinction between value, technical execution and closure is central: downstream contracts need to know not merely that a transaction has appeared on a ledger, but that the relevant obligation has been discharged in a form that can be safely treated as a settled fact (Sansom & Steigerwald, 2026). Settlement-finality reform, Token Operator rules and smart-contract interfaces therefore need to be designed together, so that the legally recognised state, the technical verification surface and downstream reliance all point to the same settled facts.

2.8.2. Native-token dependency as a first-order payments risk

Gas is an execution dependency, not just a “cost”

A missing lens that should be made explicit is the native-token (“gas”) dependency. Even if a sterling stablecoin is fully backed and ultimately settles in central bank reserves, the core functions required to operate sterling money on a public permissionless ledger—transfer, mint, burn/redemption, contract administration (pause/freeze/upgrade), and emergency controls—typically require paying fees in the ledger’s native token.

This creates a structural dependency outside sterling: regulated actors (issuer, wallet providers, on/off-ramps, and potentially banks providing corridor services) must hold, custody and source a non-sterling token whose sterling price is floating. Put differently: sterling finality remains in central bank money, but payments finality “on the rail” becomes conditional on continuous access to a non-sterling execution asset.

In operational terms, blockspace can become the “apex constraint” on sterling payments: even central-bank-backed money is only as usable as the non-sterling execution market that admits transactions under stress—raising a monetary-sovereignty question about who governs “cash-now” settlement conditions.

Why this matters for convertibility and finality

This is not merely a fee issue. In stress, fee markets can reprice endogenously: congestion rises, fees spike, inclusion becomes uncertain, and prioritisation can become adversarial. If native-token availability is impaired (or fees become prohibitively expensive/volatile), mint/burn and redemption operations can be delayed even when backing assets are pristine. That can reproduce the Bank’s instability mechanism through a different channel: delay → secondary-market discount → confidence loss / run dynamics, now driven by network economics and execution constraints rather than backing-asset credit risk.

Network fees are not just usage costs: when

redemption/conversion depends on transaction inclusion capacity, congestion can make convertibility state-contingent, creating an effective (implicit) redemption haircut at the boundary in stress (see also 2.9.1.6).

2.8.3. Financial crime and sanctions perimeter: validator fee flows are not trivial

“KYC-ed users” does not remove protocol-level counterparty risk

In a permissionless deployment, each transaction mechanically transfers value to block producers/validators (and related infrastructure participants) via fee payments. Those recipients may be pseudonymous, offshore, and not subject to UK regulatory accountability. This creates a non-trivial perimeter question: even if every stablecoin user is onboarded through regulated wallets with robust AML/CTF controls, the rail continuously routes fee income to actors whose identity, location, and sanctions status may be unknown or not screenable in real time.

In a systemic payments context, that is more than “the cost of using the network”. It raises a concrete compliance question for regulated entities: how to meet sanctions obligations that prohibit making funds or economic resources available to designated persons when the protocol-level fee recipient may be indeterminate at the point of execution.

Ecosystem-subsidy externality

There is also a broader externality. Large-scale adoption of a sterling stablecoin would increase transaction volumes and therefore increase aggregate demand for the native token and fee revenue flowing through the ledger. That can strengthen the liquidity and economic viability of the wider ecosystem (including activity unrelated to the regulated stablecoin). If parts of that ecosystem facilitate illicit finance (or if some validators/infrastructure providers are implicated), a systemic sterling stablecoin risks indirectly supporting that activity through fee flows and network effects, even if the stablecoin’s own user base is tightly controlled.

Implication for the Bank’s framework

If permissionless deployments are contemplated, the Bank should treat “gas payment counterparties” (or equivalent) as part of the systemic risk assessment, not an operational footnote. At minimum, the regime should be explicit about:

- whether and how issuers/wallets are expected to manage **sanctions/AML risk arising from protocol-level fee payments**;
- what residual risk is accepted as unavoidable in a truly permissionless setting; and
- what design constraints (e.g., **accountable**

validation arrangements, accountable sequencing layers, or other governance structures) are required for a systemic sterling instrument.

2.8.4. Singleness and “garden dynamics” under congestion

The gas dependency has direct implications for singleness and interoperability. Boundary crossing (coin↔deposit and cross-issuer conversion) typically concentrates the most gas-sensitive and “must-clear” steps—mint/burn, compliance checks, and payout orchestration. Under congestion, corridor capacity is therefore more likely to bind, even if peer-to-peer transfers still “work” for those willing to pay. If boundary crossing is bottlenecked by gas availability/cost while on-rail transfers remain comparatively fast, users will rationally avoid the boundary. Those frictions can be experienced as exchange-rate-like within sterling (“stay inside the garden”), weakening practical singleness even where the legal redemption right is unchanged.

2.8.5. Cross-sector spillovers and coupling on shared public ledgers

A further risk—closely related to resilience and finality—is **cross-sector coupling created by shared blockspace**. A systemic sterling stablecoin on a public chain would not run on a dedicated payments rail; it would compete for inclusion and execution capacity alongside unrelated activity (e.g., DeFi liquidations, speculative trading bursts, bridge incidents, governance events, or client faults).

This matters because the binding constraint in stress is often not the credit quality of backing assets but **“execute-now” capacity**: congestion, fee spikes, validator disruption, or client failures driven by unrelated sectors can directly delay or raise the cost of core money functions (transfer, mint, burn/redemption, and administrative controls). The result is an endogenous instability mechanism: if routine payment journeys and (critically) coin↔deposit conversion become state-contingent on volatile fee markets, then **delays and uncertain execution can translate into discounts and confidence effects even where backing assets are pristine**.

The Bank’s evaluation framework should therefore consider not only steady-state throughput, but whether proposed deployments meaningfully **isolate payments-critical functions** from unrelated on-chain activity—through reserved capacity, deterministic inclusion policies, credible fallbacks, and clear operating assumptions—or whether they deliberately accept coupling. If coupling is accepted, the regime should be explicit that “payments-grade” service levels cannot be guaranteed under

congestion and that convertibility performance is contingent on external network conditions.

2.8.6. Supervisory expectations if permissionless deployments are contemplated

If permissionless deployments are contemplated, the Bank should treat **gas liquidity and fee-market stress as a first-order systemic dependency** within its resilience and convertibility assessment. At minimum this implies:

- credible, testable plans to maintain **mint/burn/redemption/admin actions** under severe congestion and fee spikes (incident playbooks and pre-positioned operational inventories);
- clarity on **who bears gas risk** (passed through vs absorbed) and how state-contingent costs affect conversion terms in stress;
- stress testing that explicitly includes **native-token price spikes, fee-market dysfunction, validator disruption/censorship, and client failures**, with measurable continuity targets (e.g., time-to-finality and time-to-redeem under stress); and
- explicit treatment of the **AML/CTF and sanctions perimeter** created by protocol-level fee flows, including what residual risk is acceptable and what governance/accountability constraints are required.

Core point: on a permissionless ledger, payments-grade par performance is not solely a function of backing assets and reserve settlement. It is also a function of continuous access to a volatile, non-sterling execution asset and the functioning of its fee market and validator ecosystem—so that dependency must be governed as part of payments-grade resilience.

Q15: From the above risks, in your opinion, which ones are most crucial, specifically in the context of public permissionless ledgers, that necessitate Bank’s focus and collaborative solutions?

2.8.7. Bank needs to be explicit about risk appetite, residual risk and accountability if permissionless deployment is contemplated

Rather than attempting a comprehensive ranking, we emphasise that the risks discussed in Q14 are **all material and mutually reinforcing** in a systemic payments context. The key practical point is that many of these are not “engineering issues” that can be fully solved; they require the Bank to be explicit about **risk appetite, residual risk, and accountability** if a permissionless deployment is contemplated.

With that caveat, we see three areas that particularly warrant Bank focus because they can become binding constraints on payments-grade performance:

- **Legal recognition, finality and downstream reliance.** The Bank should require clarity on whether

raw public-chain state is legally determinative, or whether an issuer or Token Operator maintains the legally recognised holder record. If the latter, the recognition rulebook becomes critical. It should define when technical events become legally recognised, when recognised events become final for downstream reliance, whether that finality can be reopened, and what happens to smart-contract, collateral or settlement chains that act on unrecognised or later-corrected balances. For systemic use, unrecognised balances should not be usable as final settlement or eligible collateral unless the risk is explicitly intermediated.

- **Accountability and incident governance.** A systemic sterling instrument needs a clearly accountable entity with the mandate and capability to act quickly in stress: pause, freeze, correct, migrate, communicate, compensate where appropriate, and co-ordinate with regulators and critical service providers. The Bank should not rely on diffuse protocol governance or informal ecosystem responses where time-critical monetary functions are affected.
- **Execution economics / native-token (“gas”) dependency:** payments continuity and convertibility performance become conditional on a **non-sterling execution asset** and a congestion-driven fee market, creating a distinct channel for delay→discount→confidence dynamics even when backing assets are pristine.
- **Perimeter and spillovers:** protocol-level fee flows to permissionless validators (sanctions/AML perimeter) and shared-blockspace coupling (unrelated activity creating congestion/finality uncertainty) can import external shocks into sterling payments operations.

Where “collaboration” helps is mainly in **making these dependencies measurable and governable**, even if not eliminable: agreeing stress scenarios and metrics (time-to-finality/time-to-redeem under congestion), shared incident playbooks and decision processes (including fork/reorg policy), and explicit expectations on how regulated entities manage (or are permitted to tolerate) sanctions/AML exposure arising from protocol-level fee flows. If such elements cannot be made credible, that itself is informative for whether systemic sterling use on a permissionless ledger is an acceptable risk.

Q16: Can you identify other risks which you believe that will have a material impact on these technologies in the future?

2.8.8. Future risks on these technologies?

Beyond the risks above, we expect several dynamics to become more important as the technology stack evolves and adoption scales including for example:

- **Layering and composability risk:** L2s, sequencers,

bridges, shared libraries and “modular stacks” can become the dominant points of failure and governance contention, even if the base chain is robust.

- **Inclusion/ordering dynamics (MEV/censorship):** incentives around transaction ordering and censorship resistance can become increasingly salient as regulated payments volumes grow, with direct implications for predictability and fairness of execution.
- **Operational concentration:** in practice, “decentralised” rails often concentrate in a small number of wallet providers, custodians, node/endpoint operators and cloud dependencies, creating common-mode operational and cyber risk.
- **Governance drift from policy requirements:** pressure to add upgradeability, pausing/freezing and administrative controls to meet regulatory objectives can change the security model and create new trust and key-management risks over time.

2.9. Question 17-18: Comments or suggestions on Section 2.3 and 2.4

Q17: Section 2.3 above outlines minor policy refinements and clarifies the details of policy positions set out in the 2023 discussion paper. As such, specific questions for feedback are not asked for each sub-section. Respondents are invited to provide general comments or suggestions on the proposals set out in this section.

2.9.1. Redemption fees

We broadly support the direction of travel in the CP: redemption frictions are not an innocuous “payments fee” choice. Excessive or state-contingent costs to exit into bank deposits can push holders toward secondary markets, allow coins to trade away from par, and amplify confidence/run dynamics. The Bank is therefore right to frame redemption fees as stability-sensitive and to state clearly that fees should not be used to deter redemption, nor to pass through losses or costs arising from asset sales to meet redemptions.

2.9.1.1. A key clarification required: payment-service fees vs redemption frictions

The CP states “*We have reviewed the responses to the discussion paper on the proposal to either prohibit or allow proportionate fees for redemptions. We are proposing not to prohibit fees, recognising that fees are an integral part of payment chains and they also exist in current payments arrangements in commercial banking.*” (Bank of England, 2025b, Section 2.3.3).

That is true for *payments services* (processing, acceptance, ancillary services), but it risks blurring a crucial distinction: **fees for a service are not the same**

thing as fees (or delays) that weaken a par redemption claim (“*I want my money back*”). For a payments-grade instrument, the policy question is not whether firms can charge for services; it is whether the sterling user can reliably exercise the par promise without encountering an exchange-rate-like toll to cross the boundary.

A helpful way to keep this sharp is an explicit taxonomy:

- **Payment-service fees:** charges for processing/acceptance/ancillary services that can exist without undermining singleness (e.g., merchant acquiring fees; wallet subscription features).
- **Redemption fees (or redemption frictions):** charges or delays applied when exercising the issuer’s par promise (coin → bank deposit). These are stability-sensitive because they can be experienced as a haircut and trigger the instability channel the CP itself points out (friction → secondary market → away-from-par → run risk).
- **Within-sterling boundary-crossing frictions:** the *all-in* price/time/reliability of moving between stablecoin and deposits (and between stablecoins via on/off-ramps). If material, they are experienced as an exchange-rate-like switching cost inside sterling, encouraging lock-in (“garden dynamics”) and weakening practical singleness.

2.9.1.2. Regulate the outcome users experience, not the label on a fee

In stablecoin payment journeys, conversion may become embedded in everyday use (merchant sweeps, payroll, treasury). Charges can be levied by wallets, sponsors, off-ramps or “partners” and labelled as transfer/processing fees even when users experience them as “**the cost of turning my £1 coin into £1 in my bank.**” If policy disciplines only “issuer redemption fees” narrowly, it risks leaving the economically relevant friction unchanged. The regime should therefore anchor requirements to the *user-experienced boundary-crossing outcome* (price + speed + reliability), not only to the fee label or the formal point in the chain where the charge is levied.

2.9.1.3. Market structure matters: fees can create gardens even if par is protected

The CP rightly emphasises the par/run channel, but there is also a market-structure channel: if on-rail transfers are cheap/free while off-rail conversion is slower/costlier/less reliable, pricing power shifts from “low-cost payments” to “lock-in at the boundary.” Even if par is legally protected, **strategic boundary pricing can create walled gardens**, concentration, and single points of failure.

Free redemption is therefore an important “outside option,” but it is not sufficient on its own. It only disciplines ecosystems if the *practical exit-and-re-*

enter route is close to neutral in (i) price, (ii) speed, (iii) reliability, and (iv) hassle—taking into account on-ramp costs, re-issuance frictions, operating hours, and failure points. If exit/re-entry is materially worse than staying inside one ecosystem, gardens persist.

2.9.1.4. Guardrails if any fees are permitted

If the Bank permits fees in any form, it should consider hard guardrails for “par in practice”:

- **Baseline route:** require at least one standard redemption route that is free or de minimis for typical users and works reliably (not a theoretical option that is slow, fragile, or operationally constrained).
- **Predictability and comparability:** permitted fees should be ex-ante disclosed, simple, and comparable across providers; capped; and preferably fixed (or tightly bounded).
- **No stress-contingent widening:** fees and terms should be non-discretionary and not widen in stress. State-contingent widening is precisely what turns “fees” into haircuts and invites secondary-market pricing.
- **No loss pass-through:** preserve the CP’s intent that fees are not used to pass through costs/losses from asset sales to meet redemptions. This should be framed as an *outcome test* (what the user pays/ receives), not just a statement of intent.

2.9.1.5. Extend discipline to the full chain, not just the issuer

If issuer redemption is “free” but the practical conversion route requires intermediaries who charge material or state-contingent fees, the user still experiences a boundary haircut. The regime may therefore need a principle along the lines of:

- No mandatory conversion path for coin↔deposit journeys may impose material, non-transparent, or stress-contingent charges on standard users, **even if the issuer’s own redemption fee is zero.**

This can be implemented as a conduct/performance expectation on the conversion corridor as a whole (including disclosure, comparability, and service-level commitments), rather than trying to police labels.

2.9.1.6. Network fees: treat as usage costs, but don’t let them become redemption haircuts

State-contingent execution costs are a generic risk for any DLT-based corridor: if fees or inclusion probability deteriorate under congestion, redemption/ conversion can acquire an implicit haircut in stress. Public permissionless ledgers (if this sort of deployment were considered) make this risk more acute because capacity allocation is fee-market mediated, shocks can arrive from unrelated on-chain activity, and (often) fees are paid in a non-sterling native token with a floating sterling price.

2.9.1.7. Anti-abuse and cost recovery: prefer operational tools to variable redemption haircuts

If the concern is abuse (e.g., churn, operational overload), the cleaner tools are operational rather than price haircuts that undermine the par promise—e.g., transparent rate limits, queueing rules, settlement windows, minimum verification/KYC standards, and resilient capacity planning. If any friction tool is used, it should be rule-based, transparent, and ex ante (not discretionary).

Q18: Section 2.4 above outlines unchanged policies from the 2023 discussion paper. As such, specific questions for feedback are not asked for each sub-section. Respondents are invited to provide general comments or suggestions on the proposals set out in this section.

2.9.2. Remuneration

We welcome the Bank’s decision to carry forward a largely unchanged baseline from the 2023 discussion paper: it helps reduce policy uncertainty and gives industry a stable target to build toward. In that spirit, we offer one set of comments on remuneration, because it is one of the few “unchanged” elements that materially shapes the monetary character of systemic stablecoins and interacts with several other design choices discussed elsewhere in this response (backing composition, corridor design, and financial-stability tools such as holding limits).

2.9.2.1. Remuneration as an anchor for “payments-grade” money

We continue to support **non-remuneration** as a sensible anchor for a payments-grade instrument. It is consistent with the policy objective of keeping systemic stablecoins oriented toward transaction and operational liquidity needs—rather than becoming a direct savings substitute—and therefore helps keep the monetary footprint closer to cash than to an investable near-money. This also aligns with the general logic that “money-like” claims should not be engineered to compete primarily on yield.

At the same time, we see that non-remuneration does not automatically eliminate cyclicity: the relative attractiveness of holding a zero-yield instrument will vary with the interest-rate environment (a point made by (Bindseil & McLellan, 2025; Ulrich Bindseil, 2025)). If users are already limiting stablecoin balances to necessary working balances, these are unlikely to be very elastic to a change in rates - though this sort of opportunity cost channel may be more material for corporates and cash managers than for households .

2.9.2.2. Why remuneration changes the regime (portfolio channel + run surface)

The Bank is right to treat remuneration as regime-defining. If systemic stablecoins become remunerated—directly or via structurally equivalent arrangements—the instrument ceases to be mainly a payments float and becomes a rate-sensitive portfolio asset competing with deposits and MMFs. That shifts demand toward being faster-moving and more interest-elastic, and (in stress) increases the flight-to-safety/ run-risk surface because the coin becomes a stronger outside option for large balances. It also raises genuine monetary-transmission questions: the spread between the return available on stablecoins (via backing yields/ passthrough) and bank deposit rates could become policy relevant and a differential rate on stablecoin vs. bank balances could be a policy lever in its own right, that shifts liquidity allocation and bank funding costs.

2.10. Question 19: Wholesale, and non-sterling (Comments or suggestions on Section 2.5)

Q19: Section 2.5 below introduces emerging policy areas that are intended to prompt further engagement with stakeholders. These areas are presented to support ongoing dialogue and to help shape the Bank’s future approach. Respondents are invited to provide general comments or suggestions on the thoughts set out in this section. Specific questions for feedback are not asked except in sub-section 2.5.3.

2.10.1. Wholesale use of systemic stablecoin

2.10.1.1. Bank anchor and baseline (RT2/ synchronisation)

We welcome the Bank’s anchor. We agree with the low risk appetite for any material shift away from wholesale settlement in central bank money, and we continue to view central bank money as the cleanest “singleness” anchor for wholesale settlement and finality.

Against that benchmark, we strongly support the **public-rail route**—in particular synchronisation/ omnibus approaches and RT2 modernisation—as the most direct way to deliver new functionality (programmability/atomic DvP/PvP, richer messaging, extended operating models).

The DSS and related experimentation programmes are well suited to test where (if anywhere) a stablecoin leg offers incremental functionality beyond what synchronised settlement in central bank money can already deliver, and to do so in a way that keeps wholesale par settlement anchored in the public rail.

2.10.1.2. What “not shifting away from CB money” means (stack placement)

It is useful to clarify what “not shifting away from settlement in central bank money” means in practice, because systemic stablecoins—while privately issued—could *reinforce* rather than undermine the role of central bank money, depending on where central bank money sits in the stack (backing, interoperability corridor, and stress-liquidity arrangements). In today’s deposit-based system, banks can expand deposits endogenously through credit creation and support very large *gross* payment volumes with relatively limited reserves, because central bank money is typically required only to settle *net* interbank flows. This “reserve efficiency” is further supported by on-us payments, netting, intraday liquidity recycling, tiering, and settlement conventions (and could rise further under tokenised-deposit architectures that make a larger share of transfers “on-us by default”, as we argued previously (WBS Gillmore Centre Financial Technology, 2024)).

By contrast, a systemic stablecoin does not expand ex nihilo through bank credit creation. Converting deposits into stablecoin is economically closer to a withdrawal into an outside-money-like instrument than a transfer within the deposit system: it reduces bank deposit funding (a **stock** incidence) and creates corresponding demand for central bank money at (i) the deposit↔coin boundary (to complete the conversion) and (ii) in the backing arrangements. Where backing is fully in central bank deposits, the footprint is therefore closer to cash: a shift from inside money (bank deposits) toward outside-money-like balances held against central bank money. Where reinvestment into securities is permitted, the “outsideness” is *moderated* rather than eliminated: since only a fraction of reserves and deposits return to banking sector, the system still acquires a distinct coin↔deposit boundary and, under hybrid backing, an additional dependence on stress-time monetisation capacity (repo/ sale/settlement plumbing) to turn gilts back into reserves at speed (see 2.1.3).

Hence the right framing is not “private money versus central bank money”, but how stablecoin with reserve access and backing impact demand for and the leverage of central bank money and monetary policy within the system:

- **as the backing anchor** (what share is held as central bank deposits);
- **as the interoperability anchor** (how coin↔deposit conversion clears in routine journeys and at peak times, including operating hours); and
- **as the stress-liquidity anchor** (what facilities or market mechanisms ensure rapid monetisation and redemption under correlated outflows).

2.10.1.3. Wholesale viability gate: credit, netting, liquidity intensity

A useful starting point is that **wholesale settlement today is not “RTGS everywhere.”** Even in systems that ultimately settle in central bank money, high-velocity wholesale activity relies on a layered settlement-liquidity stack that economises on scarce settlement balances. That stack includes implicit credit mechanisms (netting, queuing and liquidity-saving mechanisms, intraday liquidity recycling, payment re-sequencing) and explicit credit (clearing/correspondent banking, central-bank intraday facilities, auto-collateralisation). The combined effect is to avoid universal pre-funding of gross flows (WBS Gillmore Centre Financial Technology, 2024, p. 40).

Systemic **stablecoin rails, by contrast, are naturally prefunded / RTGS-like:** settlement is typically “pay-in-coin if you have coin”. Thus absent an explicit design to recreate netting/LSMs/credit, migrating high-velocity wholesale flows onto stablecoin settlement risks (i) higher intraday liquidity intensity (more prefunding for the same throughput) and (ii) liquidity fragmentation across multiple rails and pools—precisely the concerns we flagged previously (WBS Gillmore Centre Financial Technology, 2024, p. 40).

This is where Ram (2025) provides a constructive “bridge” by accepting the diagnosis but making an optimistic call to “rebuilding better”: not “RTGS everywhere,” but (i) **high-frequency intraday netting plus advanced orchestration/LSMs**, and (ii) an explicit, competitive “**credit-for-settlement**” layer supplying settlement liquidity on demand. Conceptually, that is coherent—and it highlights the core point: a stablecoin cannot itself provide system-wide elasticity. Either it is prefunded in central bank money, or it relies on **collateral monetisation and credit** (explicitly or implicitly) to avoid full prefunding. There is no third option.

This connects directly to **banking-system incidence**. Large-scale wholesale adoption can shift funding out of bank deposits into stablecoin holdings (a *stock* incidence). Yet the system’s time-critical settlement elasticity—intraday credit, market-making liquidity, and liquidity transformation in stress—must still come from banks’ balance sheets and/or ultimately central bank facilities (a *flow* function). That creates a design tension: the system may **reduce banks’ stable funding** while still **relying on banks (and/or the central bank)** to supply intraday liquidity and absorb stress at the boundary. If that is the intended endpoint, it needs to be explicit: *who* supplies elasticity, *on what terms* (limits, collateral, pricing, operating hours), and *who bears tail risk* when many actors want cash-now at once.

An explicit “credit-for-settlement” layer (of the sort proposed by Ram (2025)), is in effect, an **intraday liquidity market**, and that can become **monetary-policy-adjacent**. Once intraday liquidity is priced at scale, the separation between “payments plumbing” and “policy rates” is harder to maintain without strong repayment/segmentation rules or a clear central-bank floor/ceiling via standing facilities (Dale & Rossi, 1996; Martin & McAndrews, 2008). This is one reason intraday liquidity is often administered or near-zero cost, and why the literature emphasises gridlock externalities and welfare-sensitive dynamics in the pricing/availability of time-critical liquidity (Bhattacharya et al., 2009; Kabadjova et al., 2023; Martin & McAndrews, 2010). The push toward 24/7 operation adds genuinely new wrinkles (multi-rail fragmentation, tokenised collateral mobility, continuous operating constraints), suggesting scope for targeted research in the current context.

Recent stress episodes reinforce why “**cash-now**” **capacity** is often the binding constraint—not simply the safety of backing assets. Margin-call liquidity shocks in March 2020, the September 2019 repo spike, and the UK gilt/LDI episode in 2022 all illustrate how quickly time-critical liquidity scarcity and collateral monetisation dynamics can dominate outcomes under correlated demands for cash. The policy direction of travel—explicit tools to supply cash against collateral to non-banks in dysfunction—underscores that liquidity backstops are central to resilience. The implication for wholesale stablecoin is straightforward: if a design depends on repo/asset monetisation (hybrid backing) or on a credit/netting overlay to avoid prefunding, then its viability hinges on **stress-time liquidity arrangements** (who supplies cash, on what terms, and at what operating hours), not only on steady-state asset quality.

In short, wholesale stablecoin viability turns less on the label “stablecoin” and more on whether the design either (i) accepts the liquidity intensity of prefunding, or (ii) explicitly rebuilds the netting-and-credit stack with clear governance, incentives, and backstops. Even if the trajectory is to “rebuild the stack better,” this raises significant questions (as above) that may benefit from research effort.

2.10.1.4. Run risk and financial stability

Run and flight-to-safety dynamics are a first-order consideration in a wholesale stablecoin context. Wholesale balances are larger, faster moving, and more operationally “time-critical” than retail.

2.10.1.4.1. Two-way runs: “flight to token” vs “flight from token”

There are two distinct stress narratives, and the regime should be robust to both:

- **Deposit/MMF → stablecoin (“flight to token”).** If a systemic stablecoin is perceived as safer, more final, or simply more usable in stressed markets (e.g., better settlement utility, fewer frictions, longer operating hours), wholesale users (corporates, FMIIs, funds) may shift rapidly out of bank deposits and/or divest MMFs into stablecoin balances. The immediate impact is **bank funding and liquidity pressure** (stock incidence), potentially concentrated at particular banks and at the conversion corridor.
- **Stablecoin → deposits/central bank money (“flight from token”).** If the stress is instead issuer/rail-centric (operational disruption, governance uncertainty, doubts about monetisation capacity under hybrid backing, or corridor fragility), outflows run in the opposite direction. The system then tests whether redemptions can clear **fast and at par** under correlated demand.

Which direction dominates is design- and scenario-dependent. A credible wholesale regime therefore shouldn’t assume a single “run direction”; it should *stress-test both*.

2.10.1.4.2. Backing mix determines where the stress lands

(A) Run into stablecoin (deposit → coin): reserve-warehouse “sink” versus securities “recycler”. If systemic stablecoins are backed predominantly by BoE deposits/reserves held inert at the **issuer**, then net inflows into stablecoin act as a liquidity sink for the banking system: the system shifts from deposit funding toward stablecoin balances, and settlement liquidity is warehoused at the issuer rather than supporting banks’ aggregate balance-sheet elasticity. At scale, this pushes the system toward a more structural need for replacement funding and/or central-bank liquidity.

If, instead, issuers are permitted/required to reinvest part of the backing in securities (e.g., short-term gilts), then inflows recycle a material fraction of liquidity back into the banking system: reserves are re-injected via payments to asset sellers, so the incidence shifts from *aggregate drainage* toward a more distributional reallocation across banks (depending on where securities sellers bank). This does not eliminate system consequences—banks still face the funding/elasticity implications of stablecoin as an outside-money-like claim—but it changes the nature of the footprint: less pure warehousing of reserves, more reshuffling of where deposits and reserves sit.

A further step would be repo-style recycling: if the stablecoin arrangement can reliably repo or otherwise lend against securities to provide time-critical settlement liquidity, then it can behave more like an

elasticity provider. But this changes the character of the regime: it turns the stablecoin complex into a liquidity intermediary (with governance, limits, and backstop questions), rather than a simple prefunded claim.

(B) Run out of stablecoin (coin → deposit/CB money): operational capacity versus monetisation capacity. The redemption direction is a different question: not “sink vs recycler” but **how quickly reserves can be mobilised to pay out at par**.

- With predominantly **reserves/BoE deposits** backing, large redemptions are less about asset liquidity and more about **corridor/operational capacity** (hours, throughput, participant dependencies, and the ability to execute payouts at scale without delay/friction).
- With **hybrid backing** (deposits + securities), heavy redemptions introduce a second channel: **monetisation capacity under stress**. The issuer must turn securities into immediately usable settlement liquidity (repo/sale + settlement plumbing + operating hours). Even “safe” assets can become hard to monetise rapidly under correlated outflows, and forced monetisation can generate market-liquidity spillovers.

2.10.1.5. Holding limits: a retail tool that does not scale cleanly to wholesale

Holding limits may be a plausible containment tool for some retail models (although we highlight a number of issues that merit serious consideration – see 2.5), but **they are far less workable in wholesale**. Wholesale balances are large, volatile, and often operationally required (settlement buffers, margin, treasury management). A binding limit either (i) breaks usability for the wholesale use case (beyond even elasticity concerns set out in 2.10.1.3), or (ii) invites exemptions so broad that the limit ceases to be a reliable macro-prudential instrument. This reinforces the point that wholesale safety needs to be delivered primarily through **design, corridor resilience, and explicit liquidity arrangements**, not through quantity caps.

2.10.1.6. A necessary choice: bounded settlement asset vs money-like wholesale cash asset

For wholesale, there are essentially two coherent safety postures:

- **Regime A: “Bounded scheme” (contain run externalities).** Keep wholesale use narrow and accept frictions: conservative access perimeter, potentially redemption windows rather than continuous on-demand convertibility, and explicit, pre-defined stress tools (with strong corridor SLAs). This contains spillovers, but constrains money-like usefulness and broader adoption.
- **Regime B: “Money-like par claim” (absorb run pressure).** If the objective is wholesale, payments-

grade par convertibility at scale, then the regime must be able to withstand *both sides* of the run problem:

- **Deposit/MMF → stablecoin flows** without destabilising bank funding/liquidity (even though wholesale balances are not sticky and holding limits are not practical); and
- **Stablecoin → deposits/CB money redemptions** while maintaining speed-at-par under correlated outflows (with the binding constraint depending on whether backing is predominantly deposits or hybrid and monetisation-dependent).

The core policy question is therefore not simply whether wholesale stablecoin is “safe in principle”, but **what endpoint is intended**—a bounded settlement asset, or a widely held wholesale cash-management instrument—and whether the corridor, monetisation pathways, and backstops are designed accordingly.

2.10.1.7. Monetary transmission (incidence channels)

See discussion in section 2.1.7. This analysis is as relevant, or more relevant, to wholesale context.

2.10.2. Non-sterling stablecoin

***Q20:** How should the Bank seek to mitigate risks from non-sterling-denominated systemic stablecoins?*

The Bank should mitigate risks from non-sterling-denominated systemic stablecoins by focusing on functional dependence rather than foreignness as such. The relevant question is not simply whether a stablecoin is foreign-denominated, foreign-issued or operated on foreign-governed infrastructure. It is whether UK-facing users, firms, platforms or markets come to depend on a non-sterling claim for payment, settlement, collateral, liquidity, treasury or market-making functions. The risks should therefore be separated into identifiable channels.

For ordinary UK domestic activity, currency-substitution risk appears comparatively low because sterling is embedded in a mutually reinforcing system of obligations, income streams, credit creation and payment liquidity. Taxes, public obligations, mortgages and other debts create sterling liabilities; wages, revenues and domestic credit relationships make sterling the natural currency for household and firm balance sheets; accounting, pricing and legal conventions reinforce sterling as the unit of account; and sterling payment and settlement infrastructure makes sterling convenient to use. This system depends on price stability, par convertibility, safe bank money and reliable payment infrastructure. The most important defence against currency substitution

is therefore a credible, liquid and resilient sterling monetary system, not a weakening of money-grade safeguards in order to promote sterling stablecoins.

Isolated cases in which crypto-native or internationally mobile wealth is used to purchase, collateralise or invest in UK assets via cryptoassets or dollar stablecoins should not be read too quickly as evidence of general UK dollarisation. Such cases may instead show that holders of crypto or dollar-stablecoin wealth are seeking to spend, collateralise or diversify from the liquidity pool in which their wealth is held. The relevant policy questions are source-of-funds verification, AML/sanctions controls, tax treatment, banking access, FX transparency and conveyancing resilience. They would become a monetary-substitution concern only if sterling-relevant markets began to price, finance, collateralise or settle systematically in non-sterling claims.

A more plausible channel is infrastructure-dependence rather than general currency substitution. Non-sterling stablecoins may become important in UK-facing activity not because UK households and firms reject sterling for ordinary domestic purposes, but because particular markets, platforms or transaction chains organise liquidity around a foreign-currency claim. This is especially relevant in cross-border payments, platform commerce, tokenised securities, collateral movement, treasury management, exchange settlement, margining, wholesale settlement and repo-like arrangements. In these settings, currency choice may be shaped less by domestic unit-of-account preference and more by market depth, composability, liquidity pools, collateral conventions, developer tooling, wallet and exchange integration, and network effects. The risk is therefore that a non-sterling stablecoin becomes the convenient settlement, liquidity or collateral layer for sterling-relevant activity, even while sterling remains dominant for wages, taxes, mortgages, accounting and ordinary domestic payments.

This distinction also matters because non-sterling activity can be both legitimate and valuable for London as a global financial centre. The policy objective should not be to discourage all non-sterling settlement or liquidity activity in the UK. A useful analogy is euro derivatives clearing after Brexit. Euro-denominated clearing remained commercially important to London and did not imply “euroisation” of the UK economy, but it raised a distinct policy concern for the EU because important euro-market liquidity, netting and margining infrastructure sat outside the EU supervisory perimeter. The EU has sought to reduce this reliance through EMIR 3 and the Active Account Requirement, but the process illustrates that liquidity pools are path-dependent: market participants do not move simply because a

jurisdiction would prefer activity to relocate. Margin efficiencies, netting benefits, operational integration and existing market depth all matter (ECB, 2024; ESMA, 2025; European Commission, 2025). The analogy for non-sterling stablecoins is that UK-facing activity need not undergo general currency substitution for dependency risk to arise. If dollar stablecoins become the deepest liquidity, settlement or collateral layer for cross-border payments, tokenised assets, platform commerce or wholesale activity, UK-facing markets may come to rely on foreign-currency infrastructure even while sterling remains the domestic unit of account.

A non-sterling stablecoin should therefore be considered systemically relevant for UK purposes not merely when it is available to UK users, but when disruption, discounting, delayed redemption, FX conversion failure or operational outage would materially impair UK-facing payments, settlement, collateral flows, market functioning, merchant activity, treasury operations or user access. Relevant indicators include the scale and concentration of UK usage, dependence of UK-facing platforms and intermediaries, use in wholesale or collateral contexts, substitutability, corridor resilience, wallet and gateway concentration, and the extent to which the instrument is used in time-critical activity.

The flow-of-funds effects should also be described carefully. UK residents acquiring a non-sterling stablecoin do not cause sterling deposits literally to “leave” the UK: a sterling bank deposit remains a liability of a UK bank. But the holder and character of that liability may change. A granular household or business deposit may be replaced by a claim held by a foreign bank, exchange, FX intermediary or other wholesale actor. The UK banking system may therefore retain aggregate funding while becoming more dependent on cross-border, concentrated or more mobile funding. At the same time, UK residents may come to hold foreign-currency claims backed by foreign bank deposits, foreign central-bank balances, foreign sovereign debt, foreign repo or foreign money-market instruments. The risk is therefore not only deposit displacement, but funding-quality transformation and dependence on foreign monetary architecture.

Mitigation should therefore preserve London’s capacity to support legitimate multi-currency activity while ensuring that UK-facing systemic functions remain visible, resilient, substitutable and accountable. The Bank should monitor UK usage, corridor flows, backing-asset composition, wallet and gateway concentration, FX conversion activity, merchant adoption, operational incidents and signs of liquidity stress. It should also require transparent FX terms, resilient UK on/off ramps, data visibility, incident-response arrangements, and

clear treatment of UK users where a non-sterling stablecoin becomes material in UK-facing activity.

This is an area where further two-country balance-sheet analysis would be valuable. Existing work on cross-border stablecoin flows and “stablecoin sandwich” arrangements shows that stablecoin issuance, use, backing and reserve investment may sit in different jurisdictions (Ulrich Bindseil, 2025). Our related stablecoin monetary-architecture work studies the domestic incidence of stablecoin issuance (Sansom et al., 2026); the natural extension is the cross-border case, where user funding, issuer backing, reserve-asset demand and FX/corridor intermediation need not sit in the same monetary jurisdiction.

A viable sterling-denominated programmable-money ecosystem is an important complement to this mitigation, but not a substitute for money-grade safeguards. The UK should not weaken the foundations that make sterling credible in order to compete with dollar or other foreign-currency stablecoins. Proportionality should instead mean calibrating requirements to actual systemic relevance, including use-case scope, concentration, corridor performance, operational readiness, backing-asset composition and redemption resilience.(ECB, 2024; ESMA, 2025; European Commission, 2025)

***Q21:** For non-sterling-denominated systemic stablecoins issued from a non-UK entity, do you think the Bank should consider deferring to the stablecoin’s home authority?*

Yes, but only on a layered and conditional basis. Deference to a home authority may be appropriate for issuer-level prudential matters, including reserve backing, safeguarding, redemption rights, reserve management, disclosure, wind-down and failure arrangements, provided the home regime delivers equivalent outcomes and there is credible supervisory co-operation, information sharing and crisis co-ordination.

However, home-authority deference should not remove the need for UK-facing requirements where risks arise locally. The Bank should distinguish issuer-level supervision from UK corridor supervision. A non-UK issuer may be soundly regulated in its home jurisdiction, while UK users and firms remain exposed to risks through wallets, exchanges, payment gateways, merchant acceptance, FX conversion, banking on/off ramps, collateral use, data visibility, incident communication and treatment of UK users in stress. A further issue is the relationship between home-authority deference, local reserve requirements and global fungibility. If a non-sterling stablecoin relies on a single global reserve pool, fungibility may be preserved

but UK users depend heavily on the home regime. If the UK requires locally held or locally governed reserve assets for UK-facing circulation, this may improve local accountability but risks creating a distinct UK tranche of the stablecoin. A hybrid model that seeks both global fungibility and differentiated local reserve rules requires clear arrangements for reserve attribution, redemption rights, safeguarding, reserve transfer, priority and loss allocation across jurisdictions.

The Bank should therefore consider home-authority deference where equivalence is credible, but retain UK-facing powers over data reporting, operational resilience, corridor performance, FX transparency, wallet and gateway standards, incident response, interoperability and treatment of UK users in failure. The appropriate model is not simple deference or full duplication, but issuer-level prudential deference combined with host-country corridor accountability.

Q22: *If so, do you agree with the factors the Bank intends to consider? Are there additional factors the Bank should consider?*

We agree that any deference to a home authority should depend on robust prudential regulation, effective safeguarding, reliable redemption rights, credible wind-down and failure arrangements, supervisory co-operation, information sharing and crisis co-ordination.

However, equivalent outcomes should not be assessed only by reference to issuer solvency or reserve backing. The Bank should also consider the UK-facing function and infrastructure role of the stablecoin. Additional factors should include: the scale and concentration of UK usage; whether the stablecoin is used for payments, settlement, collateral, treasury, margining, market-making or as a liquidity pool; the currency, jurisdiction and asset composition of the backing portfolio; the resilience and concentration of UK wallets, gateways, exchanges, payment processors and banking on/off ramps; the fairness and transparency of FX conversion terms; and whether UK users have reliable access to redemption or conversion in stress.

The Bank should also consider whether any UK-local reserve, custody or safeguarding requirement would create a distinct UK tranche of the stablecoin, and if so how fungibility, redemption priority, reserve transfer and loss allocation would work across jurisdictions.

The relevant test should be practical rather than formal. A non-sterling systemic stablecoin may be well regulated in its home jurisdiction, but still create UK-specific risks if UK users depend on it through

fragile corridors, opaque FX conversion arrangements, concentrated gateway providers, weak local accountability or non-interoperable platform structures. Equivalent home regulation should therefore be a necessary condition for deference, not a substitute for UK-facing corridor accountability.

References

Alexander, P., Fakhoury, R., Horn, T., Panjwani, W., & Roberts-Sklar, M. (2023). Financial stability buy/sell tools: a gilt market case study. *BoE Quarterly Bulletin* 2023.

Bank of England. (2021). *Assessing the resilience of market-based finance*.

Bank of England. (2022). *Financial Stability Report - December 2022*.

Bank of England. (2025a). *Bank of England opens New Contingent Non-Bank Lending Facility for applications*. <https://www.bankofengland.co.uk/news/2025/january/boe-open-new-contingent-non-bank-lending-facility-for-applications>

Bank of England. (2025b). *Proposed regulatory regime for sterling-denominated systemic stablecoins - Consultation Paper*.

BCBS-CPMI-IOSCO. (2021). *Consultative report - Review of margining practices* (Number October).

Bhattacharya, J., Haslag, J. H., & Martin, A. (2009). Why does overnight liquidity cost more than intraday liquidity? *Journal of Economic Dynamics and Control*, 6(33), 1236–1246.

Bindseil, U., & McLellan, L. (2025, December 8). *Banning stablecoin remuneration will not protect banks’ deposits*. OMFIF. <https://www.omfif.org/2025/12/banning-stablecoin-remuneration-will-not-protect-banks-deposits/>

Dale, S., & Rossi, M. (1996). *A Market for Intra-day Funds: Does it Have Implications for Monetary Policy?* (46; BoE Working Paper).

Federal Reserve Bank of New York. (2026, February 10). *Monetary Policy Implementation: Repo and Reverse Repo Agreements*.

FSB. (2020). *Holistic Review of the March Market Turmoil* (Number November).

FSB. (2025). *Enhancing the Resilience of Nonbank Financial Intermediation*.

FSB. (2026). *Vulnerabilities in Government Bond-backed Repo Markets*. <https://www.fsb.org/uploads/P040226.pdf>

Gabor, P. (2023). *An anatomy of the 2022 gilt market crisis* (1,019; BoE Working Paper, Number 1).

Giuzio, M., Ianiro, A., Lillo, F., Macchiati, V., Sowiński, A., & Telesca, E. (2024). Assessing the liquidity preparedness of investment funds to meet margin calls in derivatives markets. *ECB Financial Stability Review*, May.

Hauser, A. (2021). *From Lender of Last Resort to Market Maker of Last Resort via the dash for cash: why central banks need new tools for dealing with market dysfunction*. 1–13.

Hempel, S. J., Isley, C., Kahn, R. J., & McCabe, P. E. (2023). Money Market Fund Repo and the ON RRP Facility. *FEDS Notes*. Washington: Board of Governors of the Federal Reserve System.

Kabadjova, A., Badev, A., Garratt, R., Heijmans, R., Kosse, A., Martin, A., Nellen, T., Nilsson, T., Pustelnikov, A., & Rivadeneyra, F. (2023). *Intraday liquidity around the world* (1089; BIS Working Papers).

Kashyap, A. (2020). *The Dash for Cash and the Liquidity Multiplier: Lessons from March 2020*.

Kumhof, M., & Salgado-moreno, M. (2024). *Quantitative easing and quantitative tightening: the money channel* (1,090; BoE Staff Working Paper).

Löber, K. (2026). *Settlement Finality in the EU: Legal Foundations, Achievements, and Future Challenges*.

Martin, A., & McAndrews, J. (2008). *Should There Be Intraday Money Markets?* (337; FRBNY Staff Paper).

Martin, A., & Mcandrews, J. (2010). A study of competing designs for a liquidity-saving mechanism. *Journal of Banking and Finance*, 34(8), 1818–1826. <https://doi.org/10.1016/j.jbankfin.2009.07.023>

Ram, R. (2025). *The Emerging Architecture of Clearing and Settlement in Digital Finance* (Gillmore Centre Insight Paper).

Ram, R. (2026a). *A Two-Layer Architecture for Digital Assets: Decoupling Asset Custody from the Ownership Record* (Gillmore Centre Insights Paper). https://warwick.ac.uk/fac/soc/wbs/subjects/finance/gillmore/publications-copy/insightpapers/two_layer_architecture_for_digital_assets/gillmore-centre-insight-paper-april-2026.pdf

Ram, R. (2026b). *From “Trust the Issuer” to “Verify the Ledger”: Designing the Token Operator* (Gillmore Centre Insights Paper). <https://warwick.ac.uk/fac/soc/wbs/subjects/finance/gillmore/publications-copy/insightpapers/designing-the-token-operator/gillmore-centre-insight-paper-2-april-2026.pdf>

Regulation of Payment Stablecoins, Pub. L. (A)(1)(A)(v), 12 U.S.C. § 5903(a)(1)(A)(v) (2026).

Rhodes, T. (2025, October 3). Letter: UK stablecoin issuer responds to BoE governor. *Financial Times Letter*.

Sansom, B. (2025a). *The Bank of England’s stablecoin regime: the missing liquidity lens and open questions on credit* (Gillmore Centre Insight Paper).

Sansom, B. (2025b, November 11). *Why liquidity, not just solvency, is the UK’s real stablecoin challenge*. Finextra. <https://www.finextra.com/the-long-read/1495/why-liquidity-not-just-solvency-is-the-uks-real-stablecoin-challenge>

Sansom, B., Skeie, D., Pantelopoulos, G., & Ram, R. (2026). *Private QE or Narrow Money? How Stablecoin Architecture Shapes Monetary Transmission*

(Forethcoming; Gillmore Centre Working Papers (Forthcoming)).

Sansom, B., & Steigerwald, R. (2026). *Closing the Past, Organising the Future: Fixed Claims and the Architecture of Safety* (Gillmore Centre Working Papers (Forthcoming)).

Ulrich Bindseil. (2025). *Regulatory responses to the financial stability implications of stablecoins* (470; SAFE Working Paper).

US Federal Reserve. (2020). *Money Market Mutual Fund Liquidity Facility*. <https://www.federalreserve.gov/newsevents/pressreleases/files/monetary20201130a2.pdf>

US SEC. (2023). *SEC Adopts Money Market Fund Reforms and Amendments to Form PF Reporting Requirements for Large Liquidity Fund Advisers*.

WBS Gillmore Centre Financial Technology. (2024). *Response to: The Bank of England’s approach to innovation in money and payments* (WBS Gillmore Centre Policy Engagement).



Gillmore Centre for Financial Technology
Warwick Business School
University of Warwick

GillmoreCentre@wbs.ac.uk