# THACH NGUYEN

## Proposal and Comment Information

**Title:** Implementing the Federal Reserve Board's Responsibilities under the GENIUS Act, R-1899

**Comment ID:** FR-2026-0033-01-C32

## Submitter Information

**Name:** Thach Nguyen

**Submitted Date:** 11/09/2026

Docket No. R-1899, RIN 7100-AH29

Please find attached my comment letter on the Board's proposed rule implementing its responsibilities under the GENIUS Act. I submit it in my individual academic capacity as a Lecturer in Finance at Sheffield University Management School, University of Sheffield. The views expressed are my own and do not represent the University of Sheffield.

The letter addresses Questions 157, 158, 172, 40, 67, 70, 66 and 174. My principal points are:

1. Operational risk calibration (Questions 157, 158). The Board's economic analysis recognises that prescriptive numerical standards may favour firms with existing scale. The issuance-based operational risk requirement, whose marginal rate falls from 2.0% to 1.0% as an issuer grows, has that effect. The economic analysis does not appear to weigh it against the competition benefits it claims, and I recommend that it do so.

2. Alternative calibrations (Question 157). A uniform rate of one to two percent removes the scale gradient. A fixed component of $1 million to $25 million steepens it. The two have opposite distributional effects and should not be combined without acknowledging this.

3. Incidence (Questions 157, 172). Evidence from US commercial banks indicates that institutions with greater market power pass the cost of liquidity requirements into margins to a lesser extent. To the extent that scale and market power go together, a declining-rate schedule adds to rather than offsets the advantage of larger issuers.

4. Remuneration prohibition and bank deposits (Question 40). Board staff analysis scales linearly with adoption. Recent euro-area evidence indicates that the effects of stablecoin adoption on bank funding and lending are nonlinear in its scale, so proportional estimates may misstate the effect at high adoption.

5. Deposit concentration limits (Questions 67, 70). The two deposit formulations address different risks. A five percent limit expressed as a share of the issuer's reserves would have constrained the exposure that mattered in March 2023; a ten percent limit expressed as a share of the bank's deposit liabilities would not. I also recommend that concentration of authorised redemption counterparties be among the risks issuers must manage.

6. Treasury market interactions (Question 66) and interest rate risk (Question 174). Liquidity stress scenarios should combine redemption pressure with impaired Treasury market conditions. I support the Board's decision not to propose a general interest-rate-risk charge, and suggest a targeted treatment of long-dated collateral in reverse repurchase agreements.

Full analysis and references are in the attached letter.

Dr Thach Nguyen
Lecturer in Finance, Sheffield University Management School, University of Sheffield

**Dr Thach Nguyen**

Lecturer in Finance

Sheffield University Management School

University of Sheffield

9 November 2026

Benjamin W. McDonough

Secretary, Board of Governors of the Federal Reserve System

20th Street and Constitution Avenue NW

Washington, DC 20551

**Re: Docket No. R-1899, RIN 7100-AH29 — Implementing the Federal Reserve Board's Responsibilities Under the GENIUS Act (91 FR 61580)**

Dear Mr McDonough,

I write in response to the Board's notice of proposed rulemaking published on 29 September 2026. I am responding in my individual academic capacity as a Lecturer in Finance at Sheffield University Management School, University of Sheffield. My research is in empirical banking, bank liquidity and financial stability, financial technology and digital assets, and my published work includes studies of bank liquidity pricing (Nguyen and Nguyen, 2022), the effects of supervisory stress testing on bank liquidity creation (Nguyen et al., 2020), and the response of stablecoins and other digital assets to monetary and interbank conditions (Nguyen et al., 2022; Nguyen et al., 2026). My publications appear under Nguyen, T. V. H., except Nguyen et al. (2026), in which I appear as Nguyen, T.

I address Questions 157, 158, 172, 40, 67, 70, 66 and 174. I have not attempted the full question set, and I have confined myself to questions on which I can offer evidence rather than general observation. The views expressed are my own and do not represent the University of Sheffield.

# Summary of principal points

- **1.** The Board states in its economic analysis that prescriptive numerical standards may favour firms with existing scale, and relies on that reasoning to justify a principles-based approach to reserve asset diversification. The proposed capital requirements run in the opposite direction: the marginal rate of the issuance-based portion of the baseline operational risk requirement declines from 2.0% to 1.0% of outstanding issuance as an issuer grows. The Board gives a rationale for that graduation, but the economic analysis does not appear to weigh its effect on entry. I recommend that it do so. **(Questions 157, 158)**
- **2.** The two alternative calibrations the Board pairs in Question 157 have opposite distributional effects. A uniform rate of one to two percent removes the scale gradient; a fixed component of $1 million to $25 million steepens it. They should not be treated as complements without acknowledging this. **(Question 157)**
- **3.** Evidence from US commercial banks indicates that the cost of regulatory liquidity requirements is passed into margins to a lesser extent by institutions with greater market power. To the extent that scale and market power go together, a declining-rate schedule adds to rather than offsets the advantage of larger issuers. **(Questions 157, 172)**

- **4.** Analysis published by Board staff on stablecoins and bank deposits scales linearly with adoption. Recent euro-area evidence indicates that the effects of stablecoin adoption on bank funding and lending are nonlinear in its scale, so estimates that scale proportionally may misstate the effect at high adoption. **(Question 40)**
- **5.** Where a quantitative concentration backstop is adopted, the two deposit formulations in Question 67 address different risks. A limit expressed as a share of the issuer's own reserves protects the issuer against the failure of a bank, and at five percent is the one that would have constrained the exposure that mattered in March 2023; a limit expressed as a share of the depository institution's deposit liabilities protects the bank against the issuer, and at ten percent would not have bound. **(Question 67)**

# Questions 157, 158 and 172: the operational risk calibration and the de novo floor

**The proposed schedule declines with scale.** The issuance-based portion of the baseline operational risk requirement would equal two percent of the first $20 billion of outstanding stablecoins, 1.5 percent of the next $30 billion, and one percent of outstanding stablecoins exceeding $50 billion.

The marginal rate above $50 billion is exactly half the rate applied to the first $20 billion. The average rate falls correspondingly: an issuer with $5 billion outstanding faces an issuance-based requirement of 2.00 percent of issuance, while an issuer with $100 billion outstanding faces 1.35 percent. The smaller issuer pays approximately 48 percent more per dollar issued.

**This sits uneasily with the Board's own reasoning on competition.** In section III.F.2.b the Board writes:

> *"The proposal's principles-based approach to reserve asset diversification may also promote competition by reducing barriers to entry. Unlike prescriptive numerical standards that may favor firms with existing scale and established counterparty relationships, the flexible framework allows new Board-supervised PPSIs to develop reserve asset strategies tailored to their specific circumstances and business models."*

The Board therefore recognises that a prescriptive numerical standard can favour firms with existing scale, and treats avoiding that effect as a benefit. The declining operational risk schedule is a prescriptive numerical standard that favours firms with existing scale by construction. So does the $5 million de novo capital floor in proposed § 247.15(b)(2), which applies for three years following approval as the higher of that floor and the risk-based requirement. So, prospectively, does alternative (ii) in Question 157.

The Board does give a basis for the graduation, and Question 158 asks about it directly. The preamble reasons that operational losses may not scale linearly with outstanding issuance because, taking unauthorised minting as the example, a perpetrator must exchange stolen stablecoins for other assets before the issuer can freeze or burn them, and the liquidity available on non-compliant venues limits how much can be exchanged. I accept that this supports a declining rate for that category of loss. It is less clear that it extends to the other operational losses an issuer faces, such as processing and settlement errors in reserve management, failures at custodians and other service providers, and legal and compliance costs, some of which may have fixed components that weigh more heavily on small issuers. The Board's calibration also draws on the operational loss experience of banking organisations that report FR Y-14 data, which the Board notes are among the largest, and that experience says little about how losses behave at the small end of the range. My submission is therefore narrower than an objection to graduation. The economic analysis claims a competition benefit from avoiding prescriptive numerical standards but does not appear to weigh the capital schedule's effect in the opposite direction, and it should.

**The two alternatives in Question 157 move in opposite directions.** Question 157 asks whether the Board should consider "(i) applying a uniform capital requirement between one and two percent of outstanding issuance value and/or (ii) additionally including a fixed capital requirement between $1 million and $25 million to account for operational losses that do not scale with outstanding issuance value".

Alternative (i) eliminates the scale gradient. Alternative (ii) steepens it: a fixed dollar requirement is, by definition, a larger proportion of a smaller issuer's balance sheet. On a $25 million fixed component, an issuer with $1 billion outstanding would bear an additional 2.5 percent of issuance, while an issuer with $100 billion

would bear an additional 0.025 percent — a hundredfold difference in proportional burden. The conjunction "and/or" invites the two to be adopted together, and if they are, the distributional effect of (ii) would offset much of what (i) achieves.

If the Board's concern is genuinely that some operational losses do not scale with issuance, I would suggest that a fixed component be paired with a reduction in the proportional rate calibrated so that the *total* requirement is approximately scale-neutral, rather than added on top of a schedule that already declines.

**Evidence on incidence.** The Board asks commenters to share data on the rate of operational losses. I do not hold operational loss data for stablecoin issuers and do not offer any. I can speak to a different question, which the calibration also turns on: who ultimately bears a prudential cost once it is imposed.

In Nguyen and Nguyen (2022) my co-author and I examine US commercial banks over 2001 to 2018 and find that higher liquidity ratios are associated with higher interest margins, consistent with an economic cost to holding liquidity. Critically, that cost is not passed through uniformly: institutions with greater market power incorporate the cost of holding liquidity into their margins to a *lesser* extent than institutions with less market power. The effective burden of a prudential requirement depends on the competitive position of the firm subject to it.

Applied here, the regulatory burden and the capacity to absorb it move in the same direction. A smaller issuer faces a higher proportional charge and has less ability to absorb it without passing it to users. The Board records that the current market is dominated by two providers, with aggregate stablecoin market capitalisation of approximately $317 billion in early 2026, and counts increased competition among the benefits of the proposal. Two points follow. First, the Board expects that Board-supervised PPSIs would generally be subsidiaries of State member banks, so the gradient will operate among bank-affiliated entrants, where a schedule that charges a small issuer twice the marginal rate of a large one favours the subsidiaries of the largest banks. Second, issuers supervised by the OCC would be subject to a different design. The OCC's proposal (91 FR 10202) sets the ongoing requirement by an evaluation of each issuer's business model and risk profile, with a $5 million de novo floor and a separate operational backstop, rather than by a formula tied to outstanding issuance, which the OCC considered but did not propose. A new entrant paying two percent of issuance under the Board's rule may therefore be competing with issuers whose requirement is not set that way, and I would encourage the Board to address that comparison in its final economic analysis.

**Recommendation.** Adopt alternative (i), a uniform rate within the one to two percent range. If a fixed component is adopted under alternative (ii), offset it with a reduction in the proportional rate so that the combined requirement does not decline in scale. If the graduated schedule is retained, explain in the final rule why the rationale given for unauthorised minting extends to operational losses generally, and address in the economic analysis the schedule's effect on entry alongside the competition benefits currently claimed.

**Question 172.** The same reasoning applies to the $5 million de novo floor. I do not object to a floor in principle — a minimum absolute loss-absorbing capacity for an untested issuer is prudent. I would suggest the Board consider whether a floor that tapers across the three-year de novo period, or that is set by reference to projected issuance rather than as a flat figure, would achieve the same prudential objective at lower cost to entry. I note that Question 69 contemplates a *longer* permitted weighted average maturity for smaller issuers (45 days, against 20 or 30 days generally), which shows the Board is willing to tailor requirements by size where it sees reason to. The cumulative effect of the capital schedule and the de novo floor on smaller issuers deserves the same consideration, in the aggregate rather than provision by provision.

# Question 40: the effect of the remuneration prohibition on bank deposits

Question 40 asks what impact a narrow prohibition on paying remuneration would have relative to a broader one, and what impact either would have on bank deposits.

**Analysis published by Board staff provides a natural starting point.** Wang (2025), a FEDS Note in which Board staff offer their own views, concludes that stablecoins may reduce, recycle or restructure bank deposits rather than simply draining them, and that the net effect depends on who demands stablecoins, what assets are converted to buy them, and how issuers manage their reserves. Using a framework in which the lending response is a function of the deposit change and a composition haircut, it estimates that each $100 billion of unrecycled deposit drain implies roughly $60 to $126 billion less lending, rising to approximately $600 billion

to $1.26 trillion in its high-adoption case, which assumes $1 trillion of stablecoin growth, no recycling of reserves to banks and issuer access to master accounts.

**That framework is linear in adoption.** Its scenario arithmetic scales proportionally and treats the money multiplier as a fixed range. Altavilla et al. (2026), using confidential granular data on euro-area banks and their borrowers, find that stablecoin adoption induces deposit substitution away from retail deposits, increases reliance on wholesale funding, and can constrain intermediation capacity. They further report that the effects on the pass-through of policy rates to bank funding costs and lending conditions are **nonlinear and depend on the scale of adoption**.

If that nonlinearity holds in the United States, the implication for Question 40 is direct. Estimates calibrated on current adoption levels may not be a reliable guide to the effect at the adoption levels the GENIUS Act framework is designed to accommodate and, if effects strengthen with scale, will understate it. A narrow prohibition that appears to have modest effects on deposits at today's scale may not have modest effects at the scale contemplated. This is an argument for the broader prohibition, to the extent that remuneration paid through affiliates or third parties would raise adoption, and for not treating evidence gathered during the early period of the regime as informative about the later period.

**A refinement on pass-through.** Wang cites evidence that more than 60 percent of funding-cost increases reach lending rates, and concludes that smaller and less digitally advanced institutions face more serious headwinds. The finding in Nguyen and Nguyen (2022) described above bears directly on this: pass-through is not uniform but varies systematically with market power. That converts an average into a distributional prediction, and it predicts by mechanism the same institutions the staff note identifies by observation. Institutions with the least market power would be expected to pass on the most of any funding-cost increase, so the incidence is likely to fall disproportionately on customers of banks with less market power, which are often smaller banks.

**A measurement point.** In Nguyen et al. (2020) my co-authors and I examine the effect of Federal Reserve stress tests between 2009 and 2016 on US bank liquidity creation. We find that liquidity creation falls on both the asset side and off-balance sheet, that the effect persists in quarters after the testing period, and that stress-tested banks *increase* liability-side liquidity creation even as asset-side creation falls. Two implications for monitoring the effects of the remuneration prohibition: a single aggregate measure of deposits or liquidity creation can net out a material reallocation and should be decomposed; and effects of a known prudential change can appear in anticipation and persist afterwards, so a monitoring window aligned only to the period of implementation will understate them.

# Questions 67 and 70: quantitative limits and the principles-based approach

Question 70 asks about the advantages and disadvantages of the proposed principles-based diversification requirement and which risks issuers should be required to measure and manage. Question 67 asks about prescribing quantitative limits instead of or in addition to that approach, and offers specific formulations.

**The case for a quantitative backstop rests on the shape of the loss function.** A principles-based standard asks the issuer to manage concentration risk prudently, exercising judgement as conditions change. That is well suited to risks where losses vary smoothly with exposure. The available evidence suggests concentration risk in this market does not behave that way.

Anadu et al. (2023), in a Federal Reserve Bank of New York staff report comparing the stablecoin runs of 2022 and 2023 with money market fund runs, estimate a discrete break-the-buck threshold of $0.99 below which redemptions accelerate. Ma, Zeng and Zhang (2026) find that run incentives arise from issuers' liquidity transformation combined with the fixed redemption price available to arbitrageurs in the primary market, and that USDC bore significant run risk because of its less concentrated arbitrage, while also being exposed to default risk in the banking sector through its concentrated deposit holdings.

Where the response to deterioration is discontinuous, judgement-based management is least reliable precisely where it matters most, because the consequences of misjudgement are not proportional to the error. That is an argument for a quantitative floor beneath a principles-based baseline, not for replacing the baseline.

**On the form such a limit should take.** Question 67 offers, among others, (i) limiting uninsured eligible deposit claims against an insured depository institution to no more than five percent of the issuer's total reserve

assets, and (ii) limiting reserves held as uninsured claims issued by a single institution to no more than ten percent of *that institution's* deposit liabilities.

These address different risks, and neither substitutes for the other. Formulation (i) is a share of the issuer's own reserves and protects the issuer against the failure of a bank. Formulation (ii) is a share of the counterparty's deposit liabilities and protects the bank against the withdrawal of the issuer. The March 2023 episode was of the first kind. Circle held approximately $3.3 billion at Silicon Valley Bank, as the Board notes. That was some 8 percent of the roughly $40 billion of reserves backing USDC (CoinDesk, 2023), but only about 2 percent of the bank's $175.4 billion of deposits at the end of 2022 (FDIC, 2023). A ten percent limit under formulation (ii) would not have bound. A five percent limit under formulation (i) would have.

I read formulation (i) as a limit on claims against each institution; if an aggregate limit on uninsured deposits is intended, the final rule should say so. On that reading, its burden on entry depends on how much of its reserves an issuer holds as uninsured deposits. The Board states that it generally believes Board-supervised PPSIs should maintain reserve assets consisting predominantly of U.S. Treasury securities. An issuer holding ten percent of its reserves as uninsured deposits would need two banking relationships under a five percent limit; only an issuer holding all of its reserves in that form would need twenty. The cost to a new entrant is therefore modest at the reserve compositions the Board contemplates, although it rises with the deposit share, and the limit could carry the same allowance for temporary concentrations arising from redemption and issuance flows that the Board already recognises. Formulation (ii) scales with the size of the bank and is unlikely to bind at a large institution; its value lies in limiting a single issuer's share of a smaller bank's funding.

**Recommendation.** Retain the principles-based baseline and add a quantitative backstop. If only one of the two deposit formulations is adopted, formulation (i) is the one that addresses the risk to stablecoin holders; formulation (ii) is a useful complement where reserves are placed with smaller institutions. Under Question 70, I would add concentration of authorised redemption counterparties to the risks issuers are required to measure and manage. Ma, Zeng and Zhang find that *more efficient arbitrage amplifies run risk* — the mechanism that supports the peg in normal conditions also transmits stress — so the number and concentration of parties able to redeem at par in the primary market is a prudential variable in its own right, and the proposal does not currently treat it as one.

I note for completeness that this recommendation and the one at Question 157 both concern scale, and are consistent. A concentration limit expressed as a share of the issuer's own reserves is proportional by construction and, at the deposit shares the Board contemplates, requires only a small number of banking relationships; a declining capital rate confers a proportional advantage that accrues to large issuers. My submission in both cases is that requirements should be neutral as to issuer scale unless there is a stated prudential reason for them not to be.

# Question 66: interactions with the market for Treasury securities

Question 66 asks how the final rule's reserve asset requirements should reflect interactions with the wider Treasury market, and what requirements would prevent disruptive effects from the management or liquidation of Treasury reserve assets.

Ahmed and Aldasoro (2025) provide the most directly relevant empirical evidence. They find that stablecoin inflows lower short-term Treasury bill yields, and that the effect is **state-dependent**, strengthening under Treasury market stress and with the scale of the stablecoin sector. The version revised in February 2026 does not estimate a separate effect for outflows, but observes that under stress there is less room to time sales, so that the impact of outflows may be larger than that of inflows. Gross and Senner (2026) model the mechanism by which redemptions deplete reserves, prompt asset sales, depress prices and erode issuer solvency, amplifying further redemptions.

Three implications for the final rule.

**First, accumulation and unwind may not be mirror images.** The Board counts increased demand for US Treasuries among the benefits of the proposal in section III.F.3. That benefit accrues during accumulation. If outflows under stress move yields by more than inflows do in normal conditions, the cost during an unwind may exceed the benefit during accumulation, and the two should not be netted in the economic analysis.

**Second, the state-dependence means the relevant scenario is a joint one.** Effects strengthen under Treasury market stress and with sector scale. Stress testing of reserve liquidation that models redemption shocks against normal market conditions will understate the binding case, which is a large redemption during a period of impaired Treasury market functioning. I would encourage the Board to specify in supervisory expectations that liquidity stress scenarios combine redemption pressure with deteriorated market conditions rather than treating them separately.

**Third, the 93-day maturity cap concentrates holdings at the shortest end.** The cap is prudent as to duration risk, but it concentrates stablecoin reserve demand in a narrow maturity segment, which is also where the price impact documented by Ahmed and Aldasoro is measured. Question 57 asks whether to permit Treasury securities with remaining maturity up to two years. From a market-impact perspective there is a case for a somewhat wider permitted maturity range, subject to a weighted-average-maturity constraint, precisely because it would distribute reserve demand across more of the curve. I recognise this trades against interest rate risk and I do not suggest the balance is obvious.

# Question 174: capital for interest rate risk

The Board considered but did not propose requiring capital equal to projected losses in the fair value of reserve assets under a scenario in which interest rates increase by two percentage points, and asks about the advantages and disadvantages of such an approach and about alternative calibrations.

**A comparative observation.** The Bank of England's June 2026 policy statement includes an interest-rate-risk component within the financial risk reserve required of systemic sterling stablecoin issuers, and requires that component to be met in central bank money (Bank of England, 2026). The component is measured under the simplified standardised approach to market risk rather than by a rate scenario of the kind the Board describes. The consequence of that design choice is that the capital requirement is coupled to reserve composition: an issuer holding a larger proportion of government securities incurs a larger interest-rate-risk component, which it must then hold in central bank money. Relief on composition is therefore partly recovered through capital, and the two cannot be calibrated independently.

I raise this not as a criticism of the UK approach but because the coupling is easy to create without intending it. If the Board adopts an interest-rate-risk charge, I would encourage it to decide deliberately whether the charge should be satisfiable in any form of capital or in a specified form, and to calibrate the charge and the reserve composition requirements jointly rather than sequentially.

**Support for the Board's decision.** The Board opted not to propose the charge because a PPSI's exposure to interest rate risk is generally minimal given the restrictions on permissible reserve assets, and I agree. The proposed 93-day maturity limit is considerably tighter than the six-month residual maturity limit set out in the United Kingdom. The duration exposure of a portfolio capped at 93 days is small, and a +200 basis point shock produces correspondingly small fair-value losses. A general interest-rate-risk charge would therefore be disproportionate to the residual exposure the maturity cap leaves. This consideration would change if the Board were to permit longer-dated securities under Question 57, which is a further reason to treat Questions 57 and 174 as a single calibration decision.

**A targeted point on collateral.** Question 174 asks what other risks the Board should consider, giving the example of a counterparty providing collateral with maturity longer than three months for reverse repurchase transactions. That exposure is not addressed by the maturity cap: proposed § 247.11(b)(5) permits reverse repurchase agreements collateralised by Treasury bills, notes or bonds, without a limit on the maturity of the collateral. The proposal addresses it in part through the credit risk requirement, which applies haircuts to collateral of 0.35 percent for residual maturities between 93 days and one year, 1.41 percent between one and five years, and 2.83 percent beyond five years (Question 151). Those haircuts are calibrated to a five-day margin period of risk for counterparty credit purposes. The rate scenario the Board considered is a different and more severe test: a two-percentage-point increase implies a fair-value loss on long-dated collateral several times the largest haircut. I would encourage the Board to consider whether a limit on the maturity of eligible collateral, or a steeper haircut schedule for longer maturities, would capture the residual interest rate exposure more precisely than a general charge. Either would complement the overcollateralisation question raised at Question 61.

# Conclusion

The Board has set out a careful and detailed framework, and my comments are directed at calibration rather than at its architecture. The central submission is that the Board has articulated a sound principle in its economic analysis — that prescriptive numerical standards can favour firms with existing scale, and that reducing barriers to entry is a benefit worth weighing — and should weigh that principle against its rationale for the graduated capital schedule, which the economic analysis does not currently appear to do.

I would be glad to provide further detail on any of the research referred to above, and to share working papers or underlying data where that would assist the Board's staff.

Yours sincerely,

**Dr Thach Nguyen**

Lecturer in Finance

Sheffield University Management School, University of Sheffield

# References

Ahmed, R. and Aldasoro, I. (2025). "Stablecoins and safe asset prices." *BIS Working Papers*, No. 1270, Bank for International Settlements, May 2025, revised February 2026.

Altavilla, C., Boucinha, M., Burlon, L., Adalid, R., Fortes, R. and Maruhn, F. (2026). "Stablecoins and monetary policy transmission." *ECB Working Paper Series*, No. 3199, European Central Bank.

Anadu, K., Azar, P. D., Cipriani, M., Eisenbach, T. M., Huang, C., Landoni, M., La Spada, G., Macchiavelli, M., Malfroy-Camine, A. and Wang, J. C. (2023). "Runs and flights to safety: Are stablecoins the new money market funds?" *Federal Reserve Bank of New York Staff Reports*, no. 1073, September 2023, revised April 2024. https://doi.org/10.59576/sr.1073

Bank of England (2026). "Sterling-denominated systemic stablecoins." *Policy statement and consultation on draft Code of Practice*, 22 June 2026.

CoinDesk (2023). "Circle Confirms $3.3B of USDC's Cash Reserves Stuck at Failed Silicon Valley Bank." *CoinDesk*, 11 March 2023.

Federal Deposit Insurance Corporation (2023). "FDIC Creates a Deposit Insurance National Bank of Santa Clara to Protect Insured Depositors of Silicon Valley Bank, Santa Clara, California." *Press Release* PR-16-2023, 10 March 2023.

Gross, M. and Senner, R. (2026). "From Par to Pressure: Liquidity, Redemptions, and Fire Sales with a Systemic Stablecoin." *IMF Working Paper* No. 26/5, International Monetary Fund.

Ma, Y., Zeng, Y. and Zhang, A. L. (2026). "Stablecoin runs and the centralization of arbitrage." *The Review of Financial Studies*, hhag072.

Nguyen, T. V. H., Ahmed, S., Chevapatrakul, T. and Onali, E. (2020). "Do stress tests affect bank liquidity creation?" *Journal of Corporate Finance*, 64, 101622.

Nguyen, T. V. H. and Nguyen, T. V. H. (2022). "How do banks price liquidity? The role of market power." *Global Finance Journal*, 53, 100736.

Nguyen, T. V. H., Nguyen, T. V. H., Nguyen, T. C., Pham, T. T. A. and Nguyen, Q. M. P. (2022). "Stablecoins versus traditional cryptocurrencies in response to interbank rates." *Finance Research Letters*, 47(B), 102744.

Nguyen, T. C., Nguyen, T. V. H., Nguyen, T. and Nguyen, B. T. (2026). “Excess Liquidity, Cryptocurrency Returns, and the Moderating Role of Economic Policy Uncertainty.” *Journal of Financial Stability*, 87, 101587.

Wang, J. J. (2025). “Banks in the Age of Stablecoins: Some Possible Implications for Deposits, Credit, and Financial Intermediation.” *FEDS Notes*, Board of Governors of the Federal Reserve System, 17 December 2025.