# Yield Restrictions and the Funding of Stablecoin Settlement

Dalin Sheng
Southwestern University of Finance and Economics
shengdl@swufe.edu.cn

July 2026

**Abstract**

Yield restrictions on stablecoin issuers move remunerated liquidity to exchanges and on-chain vaults whose claims are backed by the restricted token. These wrappers fund liquidity by redeeming tokens, which retires issuer reserves and liabilities before stress, or by placing tokens with outside investors. Worse recovery beliefs and tighter restrictions redirect funding toward redemption. With cash and claims fixed, this feedback can create safe and fragile equilibria, and sufficient netting restores uniqueness without closing either route. In a certified fully endogenous economy, routing and claim creation generate multiplicity, and funding flight alone makes the restriction worsen recovery at the safe equilibrium.

*JEL classification:* E42, G01, G21, G23, G28

*Keywords:* stablecoins; runs; redemption; settlement finality; regulatory perimeter; yield restrictions

# I. Introduction

When Regulation Q ceilings repeatedly bound in the 1970s, demand for remunerated liquidity did not disappear. Money market mutual funds expanded as competitors to capped bank deposits by offering market-rate returns on portfolios of short-term instruments. Their rapid growth helped prompt Congress's six-year phaseout of the ceilings on time and savings deposits.[1] Stablecoin legislation now repeats the experiment. Congress enacted the GENIUS Act in July 2025. Upon its statutory effective date, the Act will prohibit permitted payment-stablecoin issuers from paying interest or yield solely in connection with holding, using, or retaining a token (United States Congress, 2025). In June 2026, the Bank of England set out policy positions and opened consultation on a draft Code of Practice that would maintain the prohibition on issuer-paid holding returns, permit payment- and activity-linked rewards, and treat certain issuer-enabled third-party holding returns as potential circumvention (Bank of England, 2026). The U.S. perimeter remains contested: the Act does not define "holder," the Congressional Research Service documents platform rewards tied to stablecoin balances (Tierno and Labonte, 2026), and the OCC proposes a rebuttable presumption for issuer arrangements with affiliates or related third parties while asking whether broader rules should cover indirect payments and rewards (Office of the Comptroller of the Currency, U.S. Department of the Treasury, 2026). The restriction attaches to a legal claim, but the economic service supplied by that claim can migrate to another balance sheet. An exchange reward program, a lending receipt, or a savings vault can hold the restricted token and issue a second demandable claim that does pay.

The historical analogy, however, breaks in one instructive place. A money market fund left the banking system; its portfolio held commercial paper and Treasury bills, without deposit claims on the banks it drained. A stablecoin *wrapper* stays inside. Its backing is the regulated token itself, so the way the wrapper finances its own liquidity promise is written directly on the issuer's balance sheet. This paper asks what that connection does: whether liquidity created outside the regulated perimeter can protect the wrapper while weakening the issuer that ultimately supplies settlement, and whether the yield restriction itself changes the answer.

Events in March 2023 illustrate the architecture the model formalizes. After Circle disclosed that \$3.3 billion of USDC reserves remained at Silicon Valley Bank, USDC fell to about \$0.88 on secondary markets on March 11, 2023 (European Systemic Risk Board, 2023;

[1] See Gilbert (1986) for the regulatory history and the distinct treatment of demand deposits.

Circle, 2023). Centralized exchanges paused one-for-one USDC–dollar conversion services over the weekend (European Systemic Risk Board, 2023; Watsky et al., 2024), while redemption requests from Circle's direct customers accumulated in a backlog (Team Circle, 2023). The peg recovered after U.S. authorities announced that all SVB depositors would be fully protected and banking rails reopened (European Systemic Risk Board, 2023; Team Circle, 2023; U.S. Department of the Treasury, Board of Governors of the Federal Reserve System, and Federal Deposit Insurance Corporation, 2023).[2] The public redemption terms draw two distinctions central to the model: who may redeem directly and what an on-chain transfer does to the redemption right (Circle, 2025). Circle's minting documentation separately describes the account debit and fiat payout (Circle, n.d.).

*The mechanism in words.* A wrapper that promises par cash on demand holds raw tokens and needs liquidity before stress. It has two ways to convert backing into cash at date 0. It can tender tokens to an issuer-access facility for redemption: the issuer debits the tokens, sends cash, and both its reserve assets and its outstanding liabilities shrink *before* any stress occurs. Or it can place tokens with outside investors who pay cash and hold the tokens as ordinary junior claims. Those investors are not passive. They receive the issuer's (restricted) payout and bear the junior loss if the issuer fails, so competition sets their participation price equal to expected junior loss net of any access option, minus the issuer payout.

Two forces now operate on the same margin. When expected recovery worsens, outside investors demand more compensation, and funding tilts toward redemption. When the yield restriction tightens, the payout those same investors receive falls, they again demand more, and funding again tilts toward redemption. At an interior funding split, holding wrapper cash and claim scale fixed, the resulting reallocation increases final issuer debit and therefore reduces the issuer's post-redemption reserve scale. The reallocation reduces franchise value when the reserve margin lost through redemption exceeds the payout obligation retired; for a tighter restriction, that loss also competes with the direct gain from the lower permitted payout. Endogenous cash and claim responses can offset or reinforce the fixed-quantity scale effect. This funding-flight channel operates upstream: the wrapper's *funding market* can run on the issuer's reserve base through completed redemption before any depositor queues.

Not every tender ends in a debit. Client flows are netted, and only a share $\nu$ of the gross facility tender becomes a final issuer debit; the remainder is financed as intermediary

[2] The episode fixes the institutional structure the model formalizes: holders of the same token had different settlement rights, the wrapper layer promised par it could suspend, and beliefs about issuer recovery repriced every route at once.

inventory and stays a junior claim. This final-debit share is the paper's architecture variable, and it enters the mechanism asymmetrically: issuer *exposure* to wrapper cash scales once with $\nu$, because $\nu$ determines how much of each funded unit retires liabilities, but the belief *feedback* scales with $\nu^2$, because $\nu$ also determines how strongly the outside quote moves the tender in the first place. Exposure therefore scales once with finality, the belief feedback twice. In the affine–quadratic embedded benchmark, sufficiently low finality makes every relative-depth design a contraction and thereby restores uniqueness without closing either funding route.

*Results.* The analysis is one mechanism in five layers. First, for any recovery conjecture we close the competitive wrapper contract: cash, claim scale, withdrawals, the funding split, and the zero-profit fee are unique (Lemma 1, Proposition 1). The conditional layer already contains a testable prediction. A worse conjecture raises the private value of precautionary cash, but the marginal unit of cash is partly funded outside, and outside investors now charge more. Cash rises only if precautionary relief beats that repricing; wrappers with deeper outside funding cut liquidity harder on bad news. Belief-sensitive funding makes liquidity procyclical exactly where its protection is most valuable.

Second, we ask whether wrapper liquidity protects or weakens the issuer (Proposition 2). One more unit of cash, holding recovery, claim scale, and policy fixed, relieves the issuer's stress load: fewer forced sales are routed to authorized redemption, and the issuer, smaller after final debit, carries smaller operating commitments. Its redemption-funded share sacrifices reserve margin, reduces franchise value when the lost reserve margin exceeds the associated reduction in payout obligations, and reduces feasible stress liquidity when issuer capacity binds; the outside-placed share parks directly accessible inventory in the perimeter. Cash stabilizes the issuer if and only if relief exceeds this balance-sheet cannibalization. Three funding statistics separate the channels: the average final-debit share governs balance-sheet levels, the marginal share governs the static cash effect, and a reallocation elasticity governs belief feedback. Because the elasticity is proportional to relative channel depth $\omega(1-\omega)$, funding *diversification* minimizes static reliance on either route while *maximizing* the speed at which adverse beliefs migrate funding into redemption. This flexibility–fragility trade-off is invisible to any model with an exogenous funding mix.

Third, we isolate funding feedback in embedded fixed-cash, fixed-claim-scale benchmarks (Theorem 1). In the quartic-congestion subclass, subject to the theorem's endpoint and interiority conditions, two unit-free statistics govern the phase: feedback leverage FL, the funding-loop gain at the balanced mix, and route tilt RT, the settlement-adjusted profitability advantage of issuer-facing funding measured in curvature units. Exactly three interior funding equilibria, a stable safe state, a stable fragile state, and an unstable middle state, exist if and only if feedback leverage exceeds one and route tilt lies inside a band whose half-width is the three-halves power of the leverage excess, so differently scaled funding markets share one cusp boundary within this subclass. Under the theorem's monotone sweep conditions, slowly raising finality carries the benchmark from a unique safe state through coexistence to a unique fragile state, and persistent convex quote revision produces fold hysteresis: the same architecture supports different recoveries depending on the direction from which it was reached. Within this benchmark, estimating both statistics requires exogenous recovery-news variation, matched tender-to-debit records, and local route-quote curvature. The cusp is therefore a cross-market restriction that aggregate issuance or chain-wide burn data cannot test. A complementary affine–quadratic benchmark gives a critical finality $\nu_c$: below it every relative-depth design is a contraction, so netting restores uniqueness without closing either funding route.

Fourth, we restore full contract endogeneity and separate the loops (Theorem 2). With cash, claims, withdrawals, routing, both funding quantities, and issuer liquidity all endogenous, we exhibit an open set of primitives on which the unrestricted contract has exactly three competitive equilibria, certified by outward-rounded interval arithmetic that also excludes every other equilibrium on the full action space. At a smooth stable fixed point, the sign of a tighter restriction is governed by the four policy-forcing terms in (72). In the certified economy, the restriction would improve recovery absent the funding channel, but policy-induced funding flight alone reverses the sign at the stable safe root. The same certified decomposition traces the unstable crossing, and hence the three-root multiplicity, to the classic routing-and-claims loop. Regulating the claim without its funding chain can therefore get the comparative static wrong even where it gets the number of equilibria right.

Fifth, welfare (Theorem 3, Proposition 3). Competitive providers of funding depth ignore that depth amplifies whichever route tilt prevails. Below a threshold finality the tilt runs outward, extra depth pulls funding away from redemption and improves recovery, and the market underprovides depth; above it the externality reverses. The correcting capacity charge is exact on a stable branch; when the route-tilt and welfare-exposure numerators remain nonzero, it diverges as that branch approaches a generic fold. This limits price-only capacity regulation and motivates architectural instruments. We characterize the robust optimum: an interior finality $\nu^R$ that trades the real settlement-credit cost of nonfinality against the recovery pressure of finality, with closed-form boundaries for full netting.

Public data anchor the mechanism's inputs. Wrappers hold a fifth to a third of USDC circulation; the payout a prohibition removes is 370 to 520 basis points; the issuer's marginal reserve dollar matures within days; and in March 2023 the discount, its collapse after the deposit guarantee, and the burn series supply the price and quantity moments of the funding mechanism (Section VII.B).

*Related literature.* The migration force descends from the Regulation Q experience and the narrow-banking debate (Gilbert, 1986; Pennacchi, 2012): a restriction on regulated money relocates the liquidity promise, as in Plantin (2015), where regulation shifts activity toward shadow banking. In our setting the shadow claim is backed by the regulated claim itself, so migration keeps the issuer inside the loop, and our margin is how the wrapper finances liquidity before stress: the route chosen upstream determines how many issuer liabilities and reserve assets are retired downstream, and demand can return through an endogenous market route. The withdrawal and franchise forces build on theory and evidence on bank runs, money funds, and deposit franchises (Diamond and Dybvig, 1983; Goldstein and Pauzner, 2005; Schmidt et al., 2016; Anadu et al., 2023; Drechsler et al., 2017; Kacperczyk and Schnabl, 2013), with one deliberate inversion: our primitive bounds shut down depositor-coordination multiplicity inside the wrapper, so any multiplicity comes from market feedback outside the depositor queue. In the certified full contract, that feedback is the routing-and-claims loop; the funding flight studied here changes the policy sign. The institutions also differ from a depositor queue: the claimant requesting cash from the wrapper and the holder legally entitled to redeem from the issuer are different agents.

Stablecoin theory studies reserve information, leverage, commitment, fire sales, redemption, yield, and arbitrage access (Lyons and Viswanath-Natraj, 2023; Ahmed et al., 2024; Gorton et al., 2026; d'Avernas et al., 2026; Li et al., 2026; Li and Mayer, 2025; Gross and Senner, 2026; Ma et al., 2025; Goel et al., 2026; Balasubramaniam, 2026; Klee et al., 2026). Closest are three strands. Payment-system theory shows how netting, settlement rules, and liquidity-saving mechanisms shape credit risk, run incentives, and bank behavior (Kahn and Roberds, 1998; Pagès and Humphrey, 2005; Galbiati and Soramäki, 2010; Parlour et al., 2022; Matta and Perotti, 2024; Capponi and Chang, 2025). Where that literature varies settlement speed and liquidity policy, our finality variable is the share of a gross liquidity tender that becomes final issuer debit, and because beliefs and policy move the funding route itself, it enters exposure once and the feedback loop twice. This asymmetry drives both the phase transition and the capacity-externality reversal. Work on claims built from stablecoins studies recursive collateral reuse and liquidation cascades (Koh and Pukthuanthong, 2026), composable tokens whose upper layers draw contingent liquidity from the base protocol in stress (Lee, 2026), and the two-way transmission of runs between DeFi and Treasury markets through a fully backed coin that pays no interest (Barrios et al., 2026). Our wrapper joins this family at a different margin: the settlement chain links the balance sheets without collateral reuse, the funding of wrapper cash is chosen before stress, and the structure is closer to the ETF creation–redemption pipeline, where authorized-participant balance-sheet and inventory frictions weaken arbitrage between fund shares and underlying bonds (Pan and Zeng, 2017). A contemporaneous theory of a tradable wrapper prices dealer–AP arbitrage and signaling (Yuan, 2026); our state variable is the route financing pre-stress cash, which makes the date-0 funding split the margin through which yield policy acts. Finally, the fold geometry connects to nonlinear amplification in macro-finance (Brunnermeier and Sannikov, 2014), while its run interpretation echoes repo-market dynamics (Gorton and Metrick, 2012). Our fold lives in the date-0 funding composition rather than in an asset price, so its test data are route quotes. CEA models how prohibiting stablecoin yield affects bank lending (Council of Economic Advisers, 2026); transaction-level evidence links stablecoin primary-market activity to partner banks' payment demand, reserve exposure, and loan shares (Lee and Tou, 2026). Platform-level evidence documents material reserve- and activity-based remuneration (Huang et al., 2026). Public filings establish that issuer–platform transfers are financially material and that legal claims are layered.

Section II presents the architecture and a balance-sheet map of the mechanism. Section III lays out the model. Section IV closes the conditional contract. Section V develops the cash wedge and the two funding flights. Section VI derives the phase diagram and its observable content. Section VII treats the full contract, netting, and the policy reversal, and grounds the mechanism's magnitudes in public data. Section VIII studies welfare and robust finality. The Internet Appendix[3] contains all proofs, the general nonlinear phase theorem, primitive boundaries, the certified constructive economy, and source and scope documentation.

# II. Institutional Architecture

## A. Layered Claims and Indirect Remuneration

The raw token is a direct claim governed by the issuer's redemption terms. A wrapper is a claim on an intermediary that holds the token or deploys it in a strategy. Custodial reward

[3]The Internet Appendix is included at the end of this submission PDF.

programs, lending receipts, savings-vault shares, and strategy tokens all fit this distinction (Carapella et al., 2026). A wrapper can appear liquid in normal times even though meeting a large cash withdrawal requires selling raw tokens or unwinding an encumbered position.

Remuneration can reach the wrapper holder along two routes. Reserve-based remuneration passes through some return generated by issuer reserves; activity-based remuneration is financed by the intermediary's own lending, trading, or other strategy. We represent the distinction with payout portability: a restriction on the issuer changes relative claim demand only to the extent that the restricted payout cannot be reproduced inside the wrapper. The distinction is explicit in policy proposals that separate issuer-paid interest from qualifying third-party activity rewards while warning that third-party arrangements can circumvent the issuer perimeter (Bank of England, 2026; Garcia Ocampo, 2025).

Circle's public filings illustrate the accounting separation between issuer income and distributor payments. They disclose costs associated with Coinbase agreements but not customer-level rewards, eligible balances, or pass-through. Internet Appendix IA.IX reports the disclosed amounts and their interpretation limits. These facts establish that a financially material intermediary layer exists.

## B. Settlement Finality and the Redemption Cascade

Issuer access and token transfer are distinct events. Circle's terms, for example, allow eligible Mint accounts to redeem directly, while non-Mint holders cannot redeem directly unless they open an eligible Mint account and therefore ordinarily rely on transfers or secondary-market providers for liquidity; an on-chain transfer assigns the redemption right but does not itself redeem the token (Circle, 2025; Watsky et al., 2024). Circle's documentation states that a direct redemption debits the institution's Mint balance and sends fiat. Its Form 10-K separately identifies "pending burns" awaiting blockchain finalization. Taken together, these disclosures imply that an account debit, fiat settlement, and on-chain burn need not be simultaneous (Circle, n.d.; Circle Internet Group, Inc., 2026). In principle, then, a wrapper can receive irreversible cash from an authorized intermediary while the issuer liability remains outstanding as intermediary inventory.

Figure 1 maps the date-0 mechanism onto the three balance sheets it connects. The wrapper issues claims at par, acquires raw tokens, and obtains cash by disposing of equal backing before stress. Some tokens are placed directly with ordinary outside investors. The rest enters an issuer-access facility: matched facility and issuer records reveal gross transfer

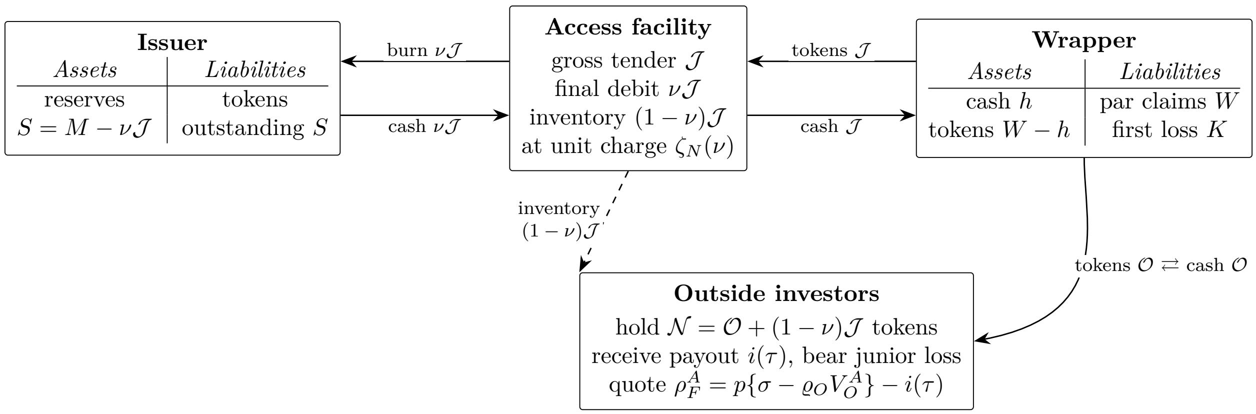


Figure 1: Date-0 funding of wrapper liquidity. Each unit of cash is funded one-for-one by gross tender $\mathcal{J}$ or outside placement $\mathcal{O}$, but only the final debit $\nu\mathcal{J}$ retires issuer reserves and liabilities; the remaining $\mathcal{N} = \mathcal{O} + (1-\nu)\mathcal{J}$ tokens stay outstanding as junior claims held in the perimeter. Outside investors receive the restricted payout, so both worse recovery beliefs (higher $\sigma$) and a tighter restriction (lower $i(\tau)$) raise their quote $\rho_F^A$ and tilt funding toward final redemption.

and completed debit. Client flows can be netted before presentation, and only completed redemption retires issuer liabilities. The paper calls the facility transfer a gross "tender" for brevity; it is not a redemption request already accepted by the issuer.

This structure separates two notions often conflated as finality. CPMI–IOSCO guidance for systemically important stablecoin arrangements defines final settlement as an irrevocable and unconditional transfer and separately discusses timely convertibility at par into other liquid assets (Committee on Payments and Market Infrastructures and Board of the International Organization of Securities Commissions, 2022). Here the wrapper–facility transfer is final, while $\nu$ caps the share that also becomes final issuer debit through completed redemption; on the implemented domain it is the realized share. The concept is distinct from blockchain confirmation and from the reversal of an accepted issuer payment. Completed redemptions cannot be clawed back at date 2; tokens retained as inventory remain junior claims. In practice $\nu$ maps onto concrete levers: same-day gross redemption versus netted settlement windows, the breadth of the authorized-participant set, and batching rules in the issuer's mint–burn pipeline. We model it as a term of the facility's regulated access agreement: the regulator sets the auditable ceiling and audit terms at date 0a, and the facility implements it because doing so is privately optimal under the sanction; finality is implemented, not commanded.

Two accounting distinctions are essential. First, value consistency requires $Px = e$:

**Date 0a: policy and contracting.** The regulator chooses the yield restriction, the issuer-debit cap $\nu$ with matched-ledger audit terms, and access capacity. Competitive wrappers post fees and first-loss terms; holders sort between wrapper and raw-token claims.

**Date 0b: noncontractible cash and balance sheets.** After fees and claim scale are fixed, wrappers choose cash and dispose of equal backing. An issuer-access facility pays cash against a gross token transfer, nets client positions, exhausts the final-debit cap when its private margin is nonnegative, complies with it when the expected sanction is large enough, and finances the residual inventory. Competitive marginal costs allocate backing between this facility and direct placement with outside raw-token investors; a per-claim reconciliation makes the facility break even. Cash is observable to the operator's claimants before date-1 requests even though it could not be contracted on at date 0a. The issuer chooses stress liquidity from the remaining reserve base.

$\Downarrow$

**Date 1: stress and final settlement.** Direct holders redeem by a cost cutoff. Wrapper holders request par cash; wrappers use cash and sell raw tokens. Ordinary buyers absorb some tokens and authorized redeemers settle the rest. Completed redemptions are final and tokens are retired.

$\Downarrow$

**Date 2: continuation or resolution.** The issuer continues or enters resolution as a function of its cash shortfall. Only tokens still outstanding receive junior recovery. Rational expectations require realized and conjectured expected haircuts to coincide.

Figure 2: Timing and settlement waterfall

below par, the tokens sold exceed the cash shortfall they finance. Second, a public burn can accompany a bridge operation, treasury management, or off-chain settlement; authorized redemption below is the economically classified route.

# III. Model

## A. Timing, Policy, and Claim Allocation

There are dates (0a), (0b), 1, and 2. Agents are risk neutral with quasilinear utility. Par is one, raw-token supply is $M > 0$, and a common stress state occurs with probability $p \in (0, 1)$. Figure 2 summarizes the order of moves.[4]

[4]Internet Appendix Tables IA.I and IA.II collect the recurring notation by model block.

The regulator chooses restriction tightness $\tau$; the permitted issuer payout is $i(\tau) = i^F - \tau \geq 0$, so $i_\tau = -1$. If $\varpi_H \in [0, 1]$ is the fraction of this payout portable through the wrapper, let $W \in (0, M)$ be aggregate wrapper claims. Marginal wrapper users have wrapper-specific convenience $\overline{A} - \mathcal{A}(W)$, where $\mathcal{A}' \geq 0$. The parameter $\varpi_H$ is a customer remuneration promise, not the share of backing that is payout eligible; write $\chi := 1 - \epsilon$ for that eligibility share. The wrapper-choice population lacks direct date-1 issuer access and is disjoint from the direct-access population introduced below, so claim clearing sorts holders on wrapper convenience alone rather than jointly on access and convenience. An operator can therefore receive an issuer payout without remitting all of it, or promise customer remuneration financed from other fee revenue. Committed settlement primitives $(\nu, q_A, f_A, \varrho_O, F_O)$ are held fixed in the conditional equilibrium and suppressed from function arguments unless varied explicitly.

## B. Two Funding Routes and the Final-Debit Share

Wrappers issue $W$ claims at par and acquire $W$ raw tokens. A symmetric operator chooses cash ratio $\lambda \in [0, \bar{\lambda}]$, with $\bar{\lambda} \in (0, 1]$, so aggregate inventory is $h = W\lambda$. Let $j \in [0, h]$ be gross facility tender and $o = h - j$ direct outside placement.

*Settlement credit.* At date 0a the contract sets an auditable issuer-debit cap $\nu \in [0, 1]$. After observing $j$ and netting client flows, the facility chooses the realized share $v \in [0, 1]$ and finances the residual inventory $(1-v)j$. Doing so consumes real resources, including inventory operations, collateral carry, and uncovered-credit loss, summarized by a unit settlement charge $\zeta_N(v)$ per token tendered, with

$$\zeta_N(v) := c_N(1-v) + \frac{g_N}{2}(1-v)^2, \qquad m_N(v) := -\zeta_N'(v) = c_N + g_N(1-v), \qquad c_N \geq 0,\ g_N > 0. \tag{1}$$

Internet Appendix Proposition IA.1(i) derives (1) from an explicit collateral problem with carry cost $\gamma_C$ and uncovered-exposure cost $\gamma_E$, giving $g_N = \gamma_C\gamma_E/(\gamma_C + \gamma_E)$; principal is a transfer within the welfare boundary. A matched-ledger audit detects $(v - \nu)_+ j$ excess debit with probability $q_A \in (0, 1]$ and assesses sanction $f_A \geq 0$ per excess token. Taking the quoted outside adjustment $\rho_F^A$ as given, strict convexity gives, for $0 < \nu < 1$,

$$v^q = \nu \iff 0 \leq \rho_F^A(\sigma, \tau) + m_N(\nu) \leq q_A f_A. \tag{2}$$

At $\nu = 0$, only the upper inequality is needed; at $\nu = 1$, only the lower. Retained inventory makes the cap bind from below; the sanction deters excess debit. Maintaining (2) on each reached state, write

$$\mathcal{J}^{\text{fin}} := \nu\mathcal{J}, \qquad \mathcal{N} := \mathcal{O} + (1-\nu)\mathcal{J} = h - \mathcal{J}^{\text{fin}}. \tag{3}$$

If either inequality fails, results must use the realized $v^q$, not the announced cap. Under implementation, issuer and perimeter cash fund $\mathcal{J}^{\text{fin}}$ and $\mathcal{N}$, the facility is flat in principal, and its tariff passes through the settlement charge.

*Outside investors.* Perimeter investors receive the issuer payout because they hold ordinary payout-eligible raw tokens. Allow a share $\varrho_O \in [0,1]$ to acquire direct issuer access at real exercise cost $c$, with cdf $F_O$ (extended by $F_O(c) = 0$ for $c < 0$) and issuer fee $\varphi_D$. Conditional on stress, their access option and exercised inventory share are

$$V_O^A(\sigma) := \int_0^{[\sigma-\varphi_D]_+} (\sigma - \varphi_D - c)\,\mathrm{d}F_O(c), \qquad a_O(\sigma) := \varrho_O F_O(\sigma - \varphi_D). \tag{4}$$

If the gross cash–token exchange is recorded at par and a perimeter investor receives signed adjustment $r_O$, competition sets

$$r_O = \rho_F^A(\sigma,\tau) := p\{\sigma - \varrho_O V_O^A(\sigma)\} - i(\tau), \qquad \rho_{F,\sigma}^A = p(1-a_O), \qquad \rho_{F,\tau}^A = 1. \tag{5}$$

The adjustment is a risk rebate when expected junior loss net of the access option exceeds the issuer payout, and a placement premium in the opposite case. This quote is where both shocks enter the funding market: worse beliefs raise it at rate $p(1-a_O)$, and a tighter restriction raises it one-for-one because the restriction lowers the payout every outside-held token earns. The full-finality, no-outside-access benchmark is $(\nu, \varrho_O) = (1, 0)$, for which $\rho_F^A = p\sigma - i(\tau)$.

*The funding problem.* Write $C_I^0(j;\kappa_F)$ for real issuer-redemption access and processing cost and $C_O^0(o;\kappa_F)$ for real outside-placement and intermediation cost, excluding the signed investor transfer; both are convex, and $\kappa_F$ collects their market-depth primitives. Competitive funding intermediaries clear through a member-owned facility described below. The private allocation chooses the unique facility tender

$$\begin{aligned}\mathcal{J}(h,\sigma,\tau;\kappa_F) := \operatorname*{arg\,min}_{0\le j\le h} \Big\{ & C_I^0(j;\kappa_F) + C_O^0(h-j;\kappa_F) + \zeta_N(\nu)j \\ & + \rho_F^A(\sigma,\tau)(h-\nu j)\Big\}, \qquad \mathcal{O} := h - \mathcal{J}.\end{aligned} \tag{6}$$

Let $C_F^P$ denote the minimized private value. Gross tender and outside placement fund cash one for one, while only final issuer debit and perimeter cash reach the two balance sheets after netting: $\mathcal{J}+\mathcal{O}=\mathcal{J}^{\text{fin}}+\mathcal{N}=h$. Since issuer payout and expected junior-loss compensation are internal transfers, integrated real funding and settlement cost is

$$C_F^R(h,\sigma,\tau;\kappa_F) := C_I^0(\mathcal{J};\kappa_F)+C_O^0(\mathcal{O};\kappa_F)+\zeta_N(\nu)\mathcal{J}. \tag{7}$$

Assume the private objective is strictly convex with the stated one-sided endpoint conditions, so the allocation is unique, and impose on every reached interior split the local strong-curvature condition

$$(C_I^0)''(\mathcal{J})+(C_O^0)''(\mathcal{O}) \geq \underline{\kappa}_F > 0. \tag{8}$$

*Three funding statistics.* At an interior split the first-order condition is $(C_I^0)'(\mathcal{J})+\zeta_N(\nu)=\nu\rho_F^A+(C_O^0)'(\mathcal{O})$, and

$$\begin{aligned}
\theta_F(h,\sigma,\tau) &:= \mathcal{J}_h = \frac{(C_O^0)''(\mathcal{O})}{(C_I^0)''(\mathcal{J})+(C_O^0)''(\mathcal{O})}, \quad \vartheta(h,\sigma,\tau) := \mathcal{J}_h^{\text{fin}} = \nu\theta_F,\\
\varphi(h,\sigma,\tau) &:= \mathcal{J}_\sigma^{\text{fin}} = \frac{\nu^2 p(1-a_O)}{(C_I^0)''(\mathcal{J})+(C_O^0)''(\mathcal{O})}, \quad \eta_F(h,\sigma,\tau) := \mathcal{J}_\tau^{\text{fin}} = \frac{\nu^2}{(C_I^0)''(\mathcal{J})+(C_O^0)''(\mathcal{O})},
\end{aligned} \tag{9}$$

with $\eta_F=\varphi/[p(1-a_O)]$ whenever $a_O<1$; all objects are evaluated at the equilibrium allocation. (Here and below, the unsubscripted $\varphi$ denotes the funding-flight elasticity; the subscripted $\varphi_D$ and $\varphi_R$ denote redemption fees.) These statistics have distinct economic meanings. The gross-tender share $\theta_F\in(0,1)$ measures how an additional unit of cash enters the facility, and the effective marginal share $\vartheta\in[0,1]$ how much of it becomes a final issuer debit. The funding-flight elasticity $\varphi\geq 0$ measures how much a *fixed* cash stock migrates into final debit when beliefs worsen, and the policy-funding elasticity $\eta_F\geq 0$ the same migration when the restriction tightens. The reason both reallocation elasticities carry $\nu^2$ while the marginal share carries $\nu$ once is the heart of the mechanism: $\nu$ enters the funding first-order condition as the weight on the outside quote. Each tendered token avoids paying the adjustment on $\nu$ units, so $\nu$ scales how strongly a repriced quote moves the tender; and it then scales again how much of the moved tender is finally debited. Finality loads once on exposure and twice on feedback. At a persistent single-channel corner the one-sided reallocation derivatives satisfy $\varphi=\eta_F=0$. The average issuer-funded share is $\bar{\omega}:=\mathcal{J}^{\text{fin}}/h$ for $h>0$ (its $h\downarrow 0$ limit at zero cash). In general $\bar{\omega}$, $\vartheta$, and $\varphi$ are three independent local statistics: they govern balance-sheet levels, the static cash effect, and belief-driven reallocation; $\eta_F$ is the policy counterpart of $\varphi$, not a fourth independent elasticity.

For the operator's cash choice, the two incidence terms that matter are

$$C^P_{F,h\sigma} = p(1-a_O)(1-\vartheta), \qquad C^P_{F,h\tau} = 1-\vartheta. \tag{10}$$

Worse recovery and a tighter yield restriction raise marginal cash cost in proportion to the marginal perimeter-funded share. Internet Appendix Lemma IA.5 gives the complete private and real envelopes.

*Decentralization.* The member-owned facility decentralizes the allocation with a marginal quote and a common per-claim reconciliation:

$$q_F(h,\sigma,\tau) := C^P_{F,h}(h,\sigma,\tau), \qquad t_F(h,W,\sigma,\tau) := \frac{C^P_F(h,\sigma,\tau) - q_F(h,\sigma,\tau)h}{W}, \tag{11}$$

$$q_F h + W t_F = C^P_F. \tag{12}$$

An atomistic claim cell pays $q_F$ on its own cash deviation and takes the reconciliation as given; at the symmetric allocation the facility breaks even and free entry capitalizes $C^P_F/W$.

*Depth and a quadratic benchmark.* For the architecture-investment results, specialize only how aggregate market depth scales route cost, not the shape of that cost. Let $d > 0$ denote proportional funding depth and write

$$\begin{gathered} C^0_I(j;d) + C^0_O(h-j;d) = \frac{\Phi_F(j;h,\omega)}{d}, \\ \Phi_{F,jj} > 0, \qquad \Phi_{F,j}(\mathcal{J}_0;h,\omega) = 0, \quad \mathcal{J}_0 := \omega h \in (0,h). \end{gathered} \tag{13}$$

The nonnegative, three-times continuously differentiable route potential $\Phi_F$ may be asymmetric away from its cost-minimizing split $\mathcal{J}_0$; assume $\Phi_F(\mathcal{J}_0;h,\omega) > 0$ in the infrastructure stage. Greater depth reduces the level and curvature of intermediation cost without imposing a linear allocation response; this is a convex proportional-depth technology, not a quadratic assumption. A quadratic implementation makes the statistics explicit: for $\omega \in (0,1)$,

$$C^0_I(j;\omega) = \frac{\kappa_0}{2\omega}j^2, \qquad C^0_O(o;\omega) = \frac{\kappa_0}{2(1-\omega)}o^2, \qquad d(\omega) := \frac{\omega(1-\omega)}{\kappa_0}, \tag{14}$$

which is the proportional-depth class with $d = 1/\kappa_0$ and joint route depth $d(\omega)$. Write $r_N(\nu,\sigma,\tau) := \nu\rho^A_F(\sigma,\tau) - \zeta_N(\nu)$ for the settlement-adjusted funding tilt. Whenever both channels are active,

$$
\begin{aligned}
\mathcal{J} &= \omega h + d(\omega)\, r_N, & \mathcal{O} &= (1-\omega)h - d(\omega)\, r_N, & & \\
\vartheta &= \nu\omega, & \varphi &= \nu^2 p(1-a_O)\, d(\omega), & \eta_F &= \nu^2 d(\omega).
\end{aligned} \tag{15}
$$

The outside channel closes at $r_N \geq \kappa_0 h/\omega$, the facility channel at $r_N \leq -\kappa_0 h/(1-\omega)$, and both reallocation elasticities are locally zero at either persistent corner. Positive net expected loss tilts funding toward final issuer debit; settlement cost and issuer payout tilt it toward perimeter placement. Internet Appendix Lemma IA.5 records the minimized private cost and the exact corner identities.

## C. Issuer Balance Sheet and Franchise

Wrapper assets satisfy $h+(W-h)=W$, while issuer reserve assets and outstanding liabilities both become $S(h,\sigma,\tau;\kappa_F) := M - \mathcal{J}^{\mathrm{fin}}(h,\sigma,\tau;\kappa_F)$. Original direct holders retain $M-W$ tokens. Outside purchasers hold $\mathcal{N}(h,\sigma,\tau)$ fully payout-eligible perimeter tokens. If a fraction $\chi = 1-\epsilon$ of remaining wrapper backing is payout eligible, the payout-eligible base is $E(W,h,\sigma,\tau;\kappa_F) := (M-W) + \mathcal{N}(h,\sigma,\tau) + \chi(W-h) = M - \epsilon W + \epsilon h - \mathcal{J}^{\mathrm{fin}}(h,\sigma,\tau) > 0$. Writing $\mu_A := r_A - c_0$ for the issuer's reserve margin, private continuation value is

$$
V^I(W,h,\sigma,\tau;\kappa_F) = \overline{V} + \xi\{\mu_A S(h,\sigma,\tau;\kappa_F) - i(\tau)E(W,h,\sigma,\tau;\kappa_F)\}, \qquad \xi > 0, \quad \epsilon \in [0,1], \tag{16}
$$

with primitive derivatives

$$
\begin{aligned}
V^I_W &= \xi i\epsilon, \qquad V^{I,0}_\tau := \xi E > 0, \qquad V^I_h = \xi[-\vartheta\mu_A + i(\vartheta - \epsilon)], \\
V^{I,F}_\sigma &:= V^I_\sigma\Big|_{W,h} = \xi(i-\mu_A)\varphi, \quad V^I_\tau\Big|_{W,h} = V^{I,0}_\tau + \xi(i-\mu_A)\eta_F.
\end{aligned} \tag{17}
$$

A tighter restriction preserves continuation value directly but also moves funding toward final redemption; the second effect has the sign of $i-\mu_A$. More wrapper cash reduces continuation value when

$$
\vartheta\mu_A > i(\tau)(\vartheta - \epsilon), \tag{18}
$$

which nests the full-redemption condition at $\vartheta = 1$. Redemption destroys reserve margin; outside placement can also expand the fully eligible payout base.

## D. Stress: Redemption, Routing, and Par-Cash Finance

Stress tests the architecture at three points: original holders decide whether to redeem, wrappers must convert backing into par cash, and every sold token must find a destination. Each point feeds the issuer's cash load.

*Direct redemption.* Let $\alpha \in (0,1]$ be the loss per outstanding junior token if the issuer enters resolution and let $\bar{\pi} \le 1$ bound the conditional resolution probability. The expected conditional haircut conjecture lies in $\sigma \in [0, \bar{\sigma}]$, where $\bar{\sigma} := \alpha\bar{\pi}$. A routine quantity $Z_0$ redeems for liquidity reasons. An issuer-access base $\bar{Z}(W)$ has heterogeneous real cost $c$ with cdf $F_D$. Direct redemption pays $1 - \varphi_D$, so a holder redeems when $c \le \sigma - \varphi_D$, and aggregate direct settlement by original holders is $\widehat{Z}(\sigma, W) := Z_0 + \bar{Z}(W)F_D(\sigma - \varphi_D)$, with $\widehat{Z}_\sigma = \bar{Z}(W)f_D(\sigma - \varphi_D) \ge 0$. This is the classic feedback channel: worse expected recovery induces settlement before resolution. Perimeter investors add direct settlement from inventory $\mathcal{N}$, so the total direct order is

$$Z^{\nu A}(\sigma, W, h) := \widehat{Z}(\sigma, W) + a_O(\sigma)\mathcal{N}(h, \sigma, \tau). \tag{19}$$

Below $Z$ denotes this total; the no-access benchmark sets $a_O = 0$. For welfare, the aggregate real access cost of an implemented original-holder quantity $Z_D \in [Z_0, Z_0 + \bar{Z}(W)]$ is $C_D(Z_D; W) := C_{D,0} + \bar{Z}(W)\int_0^{F_D^{-1}((Z_D - Z_0)/\bar{Z}(W))} c\,\mathrm{d}F_D(c)$, where redemption fees are transfers and are excluded.

*Junior claims.* A competitive continuum of infinitesimal wrapper operators supplies aggregate claims $W$; free entry makes the mass of active claim cells endogenous. In a symmetric profile, aggregate liquid inventory is $h$, early requests are $Wm$, and token sales are $x$. Let $Z$ and $Q$ denote completed direct and authorized redemptions. Because secondary sales transfer tokens but do not retire them, the stock of junior issuer claims after date-1 settlement is $N_J = S(h, \sigma, \tau; \kappa_F) - Z - Q = M - \mathcal{J}^{\mathrm{fin}}(h, \sigma, \tau; \kappa_F) - Z - Q$. Only completed redemption retires issuer exposure; legal segregation of the wrapper does not make its raw-token backing bankruptcy remote from the issuer.

*Routing.* For a wrapper sale $x$, ordinary buyers retain $q_B$ tokens and authorized redeemers settle $Q$, so $q_B + Q = x$. Ordinary buyers have aggregate expected value $(1 - \sigma)q_B - A_B(q_B)$, with strictly increasing marginal absorption cost $a_B = A_B'$, and inverse bid $P_B = 1 - \sigma - a_B(q_B)$. Authorized redeemers have capacity $\bar{Q}$, real processing cost $C_R(Q)$, and issuer fee $\varphi_R$.

Conditional on $(x, \sigma)$, the competitive route is the unique solution

$$Q^M(x, \sigma) = \arg \min_{0 \le Q \le \min\{x, \bar{Q}\}} \{A_B(x - Q) + C_R(Q) + (\varphi_R - \sigma)Q\}. \tag{20}$$

Normalize $A_B(0) = C_R(0) = a_B(0) = 0$; on the reached domain $a'_B > 0$, $c_R \ge 0$, $c'_R > 0$, and $\bar{\sigma} + a_B(\bar{x}) < 1$. The transaction price is set by the marginal destination of a sold token:

$$P^M(x, \sigma) = \begin{cases} 1 - \sigma - a_B(q_B^M), & q_B^M > 0, \\ 1 - \varphi_R - c_R(x), & Q^M = x < \bar{Q}, \end{cases} \tag{21}$$

with $P^M(0, \sigma) = 1 - \sigma$; all equilibrium allocations in the certified economy have $q_B^M > 0$, and the reached domain excludes the nongeneric simultaneous boundary $Q^M = x = \bar{Q}$. At an interior route, $a_B(x - Q) + \sigma = \varphi_R + c_R(Q)$, so $Q_x^M = a'_B/(a'_B + c'_R) \in (0, 1)$ and $Q_\sigma^M\big|_x = 1/(a'_B + c'_R) > 0$. A larger expected haircut makes priority settlement more valuable, so stress routes more of any sale to the issuer.

*Par-cash finance.* Sale proceeds are $\mathcal{B}(x, \sigma) := P^M(x, \sigma)x$, with $0 < \underline{\beta} \le \mathcal{B}_x \le 1$ and $\mathcal{B}_\sigma \le 0$ on the reached domain. If early requests exceed cash, the wrapper's par-cash shortfall and financing sale are

$$e = [Wm - h]_+, \qquad x = \Xi(e, \sigma) := \mathcal{B}(\cdot, \sigma)^{-1}(e). \tag{22}$$

Feasible par service requires $\mathcal{B}(W - h, \sigma) \ge [Wm - h]_+$ at every reached candidate state. At a smooth active point, $\Xi_e = 1/\mathcal{B}_x > 0$ and $\Xi_\sigma = -\mathcal{B}_\sigma/\mathcal{B}_x \ge 0$: a worse conjecture requires weakly more tokens to finance the same cash.

## E. Withdrawals, First Loss, and the Operator's Cash Choice

The wrapper side of the mechanism is a standard demandable-claim problem with one twist: the marginal cost of holding cash is priced by the funding market of Section III.B, so it moves with beliefs and with policy. We set up the claim waterfall first and then the operator's choice.

*Loss waterfall.* Let $Y := W - h - x$ be retained raw-token backing and write per-claim stocks as $s := x/W$ and $y := Y/W = 1 - \lambda - s$. Let $\phi(0) = 0$ and $\phi' \ge 0$ denote conversion cost per claim, so aggregate conversion cost is $W\phi(s)$. Per-claim loss if the issuer continues and if it enters resolution is $\ell_0 = \phi(s) + (1 - P)s$ and $\ell_1 = \ell_0 + \alpha y$. First-loss capital per claim is

$K \geq 0$ with cdf $H_K$, common knowledge at date 0a but realized after date-1 requests and sales; holders therefore use expected residual loss. Shocks are independent across the continuum of claim cells and independent of issuer resolution conditional on the common stress state, so aggregate incidence is deterministic; the capital is outside operator wealth pledged against losses and neither funds date-1 cash nor counts as backing. Define $b_K(L) := \mathbb{E}[\min\{K, L\}]$ and $\psi_K(L) := \mathbb{E}[(L-K)_+]$. With conjectured conditional resolution probability $r_\sigma = \sigma/\alpha$, expected per-claim operator and late-holder losses in stress are $b = (1-r_\sigma)b_K(\ell_0) + r_\sigma b_K(\ell_1)$ and $\psi = (1-r_\sigma)\psi_K(\ell_0) + r_\sigma \psi_K(\ell_1)$, with aggregates $B_W = Wb$ and $\Psi_W = W\psi$. Because $b_K(L) + \psi_K(L) = L$, state losses are allocated before expectations and

$$b + \psi = \ell_0 + \sigma y. \tag{23}$$

*Withdrawal cutoff.* A mass $m_0$ of claims has an immediate liquidity need and withdraws in stress; a mass $1 - \bar{m} > 0$ is contractually locked until date 2. Each holder in the remaining mass $\Delta_m := \bar{m} - m_0$ draws a private idiosyncratic exit cost $u \geq 0$ with cdf $F_H$, density $f_H$, and $F_H(0) = 0$. A flexible holder who remains bears expected dilution $D$; one who exits receives par and pays $u$. Exit is optimal exactly when $u \leq D$, so the aggregate early share solves

$$m = m_0 + \Delta_m F_H(D) =: G_H(D), \qquad D = \frac{\psi}{1-m}. \tag{24}$$

The ex ante real exit cost per wrapper claim generated by this cutoff is $t_H(D) := \Delta_m \int_0^D u \, \mathrm{d}F_H(u)$, with $t'_H(D) = \Delta_m D f_H(D) \geq 0$. For the aggregate and operator-specific problems, $j \in \{A, o\}$, we impose the primitive feedback bound

$$\kappa_j^H := \sup_{\mathcal{D}_j} \Delta_m f_H(D^j) \left| \frac{\psi_m^j}{1-m} + \frac{\psi^j}{(1-m)^2} \right| < 1, \tag{25}$$

over the maintained aggregate and fixed-price claim-cell domains $\mathcal{D}_A$ and $\mathcal{D}_o$. This modeling choice rules out an unexamined high-withdrawal root and makes both withdrawal fixed points globally unique, so multiplicity, when it occurs below, comes from system feedback outside the depositor queue.

*The operator's problem.* For a candidate common cash ratio, let $z(\lambda; W, \sigma, \bar{Q}) := (m(\lambda), P(\lambda))$ be the symmetric state generated by withdrawal, par-cash finance, routing, and market clearing at $h = W\lambda$. Price taking does not make an operator's own customers passive: given aggregate price $P$, conjecture $\sigma$, and a deviating cash ratio $\lambda^o$, its claim-cell withdrawal share solves $m^o = G_H(D^o)$, $D^o = \psi(\lambda^o, m^o; P, \sigma)/(1 - m^o)$, with unique solution $m^o = \mu(\lambda^o; P, \sigma)$ under the operator-level bound in (25); at a symmetric profile $m = \mu(\lambda; P(\lambda), \sigma)$. Internet Appendix Lemma IA.3 shows $\mu_\lambda, \mu_P \le 0$. Define fixed-market incidence derivatives $b^o_\lambda := b_\lambda|_{m,P} + b_m\mu_\lambda$ and $\psi^o_\lambda := \psi_\lambda|_{m,P} + \psi_m\mu_\lambda$. The cash ratio is not contractible at date 0a, and neither the fee nor the first-loss rule can be contingent on the date-0b realization. Let $a_0(\lambda, W)$ be per-claim installation, custody, and capital-carry cost excluding payout and funding cash flows. Given the competitive two-part tariff, a claim cell's private nonpayout cost is $a^o(\lambda^o; \lambda, W, \sigma, \tau) := a_0(\lambda^o, W) + q_F(W\lambda, \sigma, \tau)\lambda^o + t_F(W\lambda, W, \sigma, \tau)$, and at the symmetric allocation, facility zero profit implies $a(\lambda, W, \sigma, \tau) := a^o(\lambda; \lambda, W, \sigma, \tau) = a_0(\lambda, W) + C^P_F(W\lambda, \sigma, \tau; \kappa_F)/W$. The own-cash derivative is $a_{0,\lambda} + q_F$; below, $a_\lambda$ denotes this derivative evaluated symmetrically. Along an interior funding allocation, (10) gives $a_{\lambda\sigma} = C^P_{F,h\sigma} = p(1 - a_O)(1 - \vartheta)$ and $a_{\lambda\tau} = C^P_{F,h\tau} = 1 - \vartheta$. These two cross-partials are how the funding market prices beliefs and policy into the operator's marginal cost of liquidity.

Let $\rho_0\lambda^o$ be cash's outside opportunity cost. The operator receives $\chi(1 - \lambda^o)i(\tau)$ on eligible retained backing and promises the holder $\varpi_H i(\tau)$. After the fee is fixed, it minimizes

$$c^P(\lambda^o; \lambda, W, \sigma, \tau, P) := a^o(\lambda^o; \lambda, W, \sigma, \tau) + \rho_0\lambda^o + [\varpi_H - \chi(1-\lambda^o)]i(\tau) + p\, b(\lambda^o, \mu(\lambda^o; P, \sigma); P, \sigma), \tag{26}$$

holding the aggregate $(\lambda, P, \sigma, W)$ and tariff fixed. The price-taking marginal condition, evaluated at $\lambda^o = \lambda$, is

$$\mathcal{F}_H(\lambda; W, \sigma, \tau) := a_\lambda + \rho_0 + \chi i(\tau) + p\, b^o_\lambda = 0. \tag{27}$$

The regular domain maintains $b^o_\lambda < 0$ and $\psi^o_\lambda \le 0$ (Internet Appendix Lemma IA.3); the symmetric best-response stability derivative is $\mathcal{D}_H := \frac{\mathrm{d}}{\mathrm{d}\lambda}\mathcal{F}_H(\lambda; W, \sigma, \tau, P(\lambda)) > 0$.

*The cash crowd-out.* How does a worse conjecture move liquidity? Two forces compete. Precautionary relief pushes cash up because cash reduces expected first-loss incidence more when losses loom. Repricing pushes it down: the marginal unit of cash is partly funded outside, and outside investors now charge more. Isolate the first force as

$$\mathcal{E}^\sigma_H := -p\frac{\mathrm{d}b^o_\lambda}{\mathrm{d}\sigma} \quad \text{at fixed } (\lambda, W), \tag{28}$$

where the conditional withdrawal and token-market state adjust but the funding term is excluded. If $a_0$ has no direct recovery dependence, then at an interior funding split

$$\mathcal{F}_{H,\sigma} = p(1 - a_O)(1 - \vartheta) - \mathcal{E}^\sigma_H, \qquad \lambda^P_\sigma = \frac{\mathcal{E}^\sigma_H - p(1 - a_O)(1 - \vartheta)}{\mathcal{D}_H}. \tag{29}$$

Worse recovery raises cash only when precautionary loss relief exceeds the risk premium on the marginal outside-funded unit. Deep outside funding can make liquidity fall precisely when its private protection becomes most valuable. The model predicts that wrappers with larger perimeter-funded shares cut liquidity more sharply on adverse recovery news.

*Fee, all-in cost, and claim clearing.* Let $\lambda^P(\sigma, W, \tau; \bar{Q})$ be the unique symmetric solution and set $h^P := W\lambda^P$. Free entry anticipates the post-issuance cash choice and sets the net fee (a negative value is a rebate) to expected operator cost:

$$f^P = a(\lambda^P, W, \sigma, \tau) + \rho_0\lambda^P + [\varpi_H - \chi(1 - \lambda^P)]i(\tau) + p\, b. \tag{30}$$

The holder pays this fee, receives $\varpi_H i(\tau)$, and bears residual loss, so the all-in incremental cost of a wrapper claim relative to a raw token is

$$\begin{aligned}\Delta_H(W, \sigma, \tau; \bar{Q}) &:= f^P + (1 - \varpi_H)i(\tau) + p\{\psi + t_H(D) - \sigma\} \\ &= a(\lambda^P, W, \sigma, \tau) + \rho_0\lambda^P + [\epsilon + \chi\lambda^P]i(\tau) + p\{\ell_0 + \sigma y + t_H(D) - \sigma\}.\end{aligned} \tag{31}$$

The second line is the competitive-capitalization identity: customer pass-through changes the fee and promised payout one-for-one and drops out of all-in cost when the net fee is flexible. Marginal-holder indifference is

$$\mathcal{A}(W) + \Delta_H(W, \sigma, \tau; \bar{Q}) = \overline{A}. \tag{32}$$

Let $w(\sigma, \tau; \bar{Q})$ denote the unique clearing scale and define $\mathcal{D}_W := \mathcal{A}'(W) + \Delta_{H,W} > 0$.

*Migration.* At fixed $(W, \sigma, \tau)$, let $\dot{\Delta}_{H,\lambda}$ be the total cash-ratio derivative of all-in cost, allowing the conditional withdrawal and token-market state to adjust. The operator condition cancels the private first-loss derivative, leaving $\dot{\Delta}_{H,\lambda} = p\psi^o_\lambda + p\, t'_H(D)\dot{D}_\lambda + p[(b + \psi)_P + (b + \psi)_m\mu_P]P_\lambda \leq 0$ on every regular active-cash piece; Internet Appendix Lemmas IA.3 and IA.4 derive the three weakly negative terms from primitives. At a smooth interior funding and cash solution, $a_{\lambda\tau} = 1 - \vartheta$ and $i_\tau = -1$ give

$$\lambda^P_\tau = \frac{\vartheta - \epsilon}{\mathcal{D}_H}, \qquad \Delta_{H,\tau} = \frac{\mathcal{N}}{W} - [\epsilon + \chi\lambda^P] + \dot{\Delta}_{H,\lambda}\lambda^P_\tau, \tag{33}$$

$$\Delta^{W,P}_\tau := -\Delta_{H,\tau} = \epsilon + \lambda^P(\bar{\omega} - \epsilon) - \dot{\Delta}_{H,\lambda}\frac{\vartheta - \epsilon}{\mathcal{D}_H}, \qquad w_\tau = \frac{\Delta^{W,P}_\tau}{\mathcal{D}_W},$$

$$w_\sigma = -\frac{\Delta_{H,\sigma}}{\mathcal{D}_W}, \qquad h^P_\tau = W\lambda^P_\tau, \qquad h^P_W = \lambda^P + W\lambda^P_W. \tag{34}$$

The policy cash response is a two-force result: lower payout on retained backing encourages cash, the higher outside-funding adjustment discourages it, and cash rises if and only if $\vartheta > \epsilon$. Claim migration is positive if and only if $\Delta_\tau^{W,P} > 0$. Flexible fees capitalize customer pass-through but do not remove the funding-incidence channel. Let $\widetilde{Q}(\sigma, W, h; \bar{Q})$ denote the authorized route generated by the unique conditional system at fixed inventory

$$\widetilde{Q}(\sigma, W, h; \bar{Q}), \qquad \widehat{Q}(\sigma, W, \tau; \bar{Q}) := \widetilde{Q}(\sigma, W, h^P(\sigma, W, \tau); \bar{Q}). \tag{35}$$

## F. Issuer Liquidity and Equilibrium

Let $\varphi_j \in [0,1]$ and $\zeta_j := 1 - \varphi_j$ be the fee and net cash paid on direct and authorized settlement, $j \in \{D, R\}$. The issuer also faces an operating cash drain $R_X(S) \geq 0$ with $R_X'(S) \geq 0$: a larger issuer has weakly larger operating commitments. Conditional issuer cash demand is $R = R_X(S) + \zeta_D Z + \zeta_R Q$. Allowing the drain to scale with $S$ matters because date-0 wrapper cash reduces both the issuer's feasible liquidity and the operating load that liquidity must cover. The post-inventory reserve base limits stress liquidity: let $\overline{L}(S)$ be weakly increasing with $0 \leq \overline{L}(S) \leq S$, define the minimum feasible shortfall $\underline{q}(R, S) := [R - \overline{L}(S)]_+$, and let the issuer choose $\ell \in [0, \min\{R, \overline{L}(S)\}]$, equivalently shortfall $q = R - \ell$, to minimize

$$\mathcal{I}(V, R, S) = \min_{q \in [\underline{q}(R,S), R]} \{g_I(R - q) + p[F_I(q) + V\pi(q)]\}. \tag{36}$$

Normalize $g_I(0) = F_I(0) = 0$; liquidity cost $g_I$ is increasing and strictly convex, emergency-finance cost $F_I$ is increasing, and conditional resolution probability $\pi(q)$ is increasing, bounded by $\bar{\pi}$, with $\pi(0) = 0$. The maintained curvature condition $F_I'' + V\pi'' \geq 0$ gives a unique shortfall $q^*(V, R, S)$. In an unconstrained interior regime, with

$$D_I = g_I'' + p(F_I'' + V\pi'') > 0, \tag{37}$$

the responses are $q_R^* = g_I''/D_I > 0$, $q_V^* = -p\pi'/D_I < 0$, and $q_S^* = 0$. At a locally binding positive liquidity capacity, $q^* = R - \overline{L}(S)$ and $(q_R^*, q_V^*, q_S^*) = (1, 0, -\overline{L}'(S))$. At persistent zero liquidity the response vector is $(1, 0, 0)$; at persistent full coverage it is $(0, 0, 0)$.

The issuer-implied conditional expected haircut is

$$\mathcal{H}(Z, Q, V, S) := \alpha\pi(q^*(V, R_X(S) + \zeta_D Z + \zeta_R Q, S)), \tag{38}$$

with smooth-point derivatives $\mathcal{H}_Z = \alpha\pi' q_R^* \zeta_D \geq 0$, $\mathcal{H}_Q = \alpha\pi' q_R^* \zeta_R \geq 0$, $\mathcal{H}_V = \alpha\pi' q_V^* \leq 0$, and $\mathcal{H}_S = \alpha\pi'\{q_R^* R_X'(S) + q_S^*\}$. The scale derivative is deliberately unsigned: a larger scale raises the operating drain but relaxes a binding liquidity capacity. In a slack interior regime, $\mathcal{H}_S = \alpha\pi' q_R^* R_X'(S) \geq 0$; at binding positive capacity, $\mathcal{H}_S = \alpha\pi'\{R_X'(S) - \overline{L}'(S)\}$. A *policy recovery equilibrium* is a conjecture $\sigma$, claim scale $W$, and the associated unique conditional market and issuer choices such that claim clearing (32) holds and

$$\sigma = \mathcal{H}(Z, Q, V, S). \tag{39}$$

ASSUMPTION 1 (Primitive regularity and feasibility): All choices lie in compact domains, all functions are continuous and finitely piecewise twice continuously differentiable, and derivatives are taken within a smooth regime (one-sided at a kink). In addition:

(a) *Funding and settlement.* At fixed $\kappa_F$, the objective in (6) is strictly convex, has the stated endpoint signs, and satisfies (8) at every reached interior split. Competitive outside pricing, final debit, residual-inventory finance, principal closure, and the two-part zero-profit tariff are exactly (5), (3)–(7), and (11)–(12).

(b) *Feasibility.* Every reached allocation satisfies $0 \leq h \leq W < M$, $E > 0$, $V^I \geq 0$, $0 \leq \widehat{Z} \leq M - W$, $Z + Q \leq S$, and $N_J \geq 0$. Completed settlements are final, cannot be submitted twice, and every par or state-contingent promise is feasible.

(c) *Conditional contract.* The access and withdrawal distributions have bounded piecewise continuous densities. The two primitive contraction bounds in (25) hold. On every active cash piece, $b_\lambda^o < 0$ and $\psi_\lambda^o \leq 0$. On every smooth cash piece, two primitive curvature margins are uniformly positive: the own-cash curvature of an atomistic claim cell's objective at fixed aggregate price and tariff, and the curvature of the symmetric conditional cash residual; the exact bounds $\underline{c}_o > 0$ and $\underline{\Delta}_H > 0$ are Internet Appendix (IA.1). Marginal conditions have upward jumps at any cash kink, the symmetric cash residual has the KKT endpoint signs on $[0, \bar{\lambda}]$, and the competitive net fee is flexible and not at a floor. Comparative-static signs that use the routing block are asserted only on pieces with $q_B > 0$; the all-priority boundary instead uses the second price in (21) and requires its own one-sided check.

(d) *Claims and committed designs.* On a compact interval $[\underline{W}, \overline{W}] \subset (0, M)$, the claim residual has opposite weak endpoint signs, at least one strict, and the convenience-demand slope $\mathcal{A}'(W)$ uniformly dominates any negative all-in-cost slope in claim scale

(Internet Appendix (IA.2)). For every committed design used in Section VIII, the design-conditioned contractions, curvature margins, slope dominance, and endpoint inequalities hold uniformly on its stated domain.

Parts (a)–(b) close balance sheets and feasibility; part (c) supplies the conditional contraction and cash-optimality arguments; part (d) closes claim scale. Issuer liquidity uniqueness follows separately from the curvature condition in (36)–(37). The assumption deliberately does *not* impose uniqueness of the outer recovery fixed point, which is the object characterized below, and the constructive family does not need the global own-cost curvature gate: its nonconvex inactive tail is handled directly by the interval certificate.

# IV. The Conditional Contract

Conditional on a recovery conjecture, the layered market has a unique outcome. Any multiplicity must therefore arise from the recovery fixed point, rather than from wrapper-run coordination, a nonconvex issuer problem, or an indeterminate funding split.

LEMMA 1 (Existence and uniqueness of the conditional layer): *Under Assumption 1, for every conjecture $\sigma$ and candidate tuple $(W, \tau, \nu, \varrho_O, \bar{Q})$, there is a unique admissible wrapper–market subequilibrium, a unique cleared date-0 funding split $(\mathcal{J}, \mathcal{O})$, a unique symmetric competitive cash ratio $\lambda^P$, and a unique zero-profit fee $f^P$. After holder sorting, claim clearing is globally unique on the maintained claim interval.*

*Proof.* Lemma IA.1 closes routing and prices, and Lemma IA.3 closes withdrawals, par-cash finance, and the conditional market state. Lemma IA.5 gives the unique funding allocation. Lemma IA.4 then gives the unique symmetric cash choice and zero-profit fee; its claim-residual argument, together with Assumption 1(d), gives global uniqueness of claim clearing on the maintained interval. □

The conditional layer shows how beliefs and policy move the contract when the recovery conjecture is held fixed. Three margins matter: how a fixed cash stock is refinanced, whether the cash stock itself rises or falls, and whether holders migrate between claims.

PROPOSITION 1 (Conditional comparative statics): *Under Assumption 1, at a smooth regular point:*

(a) Funding flight. *At fixed cash, a worse recovery conjecture weakly raises the withdrawal share* $m$*, dilution* $D$*, the financing sale* $x$*, and authorized settlement* $Q$*, weakly lowers price* $P$*, and at an interior date-0 funding split shifts backing from perimeter inventory to final issuer debit:* $\mathcal{J}_\sigma^{\mathrm{fin}} = \varphi \geq 0$ *and* $\mathcal{N}_\sigma = -\varphi \leq 0$*, both strict at an interior split when* $\nu > 0$ *and* $a_O < 1$*. In a smooth active-cash, interior-routing regime,* $m_\sigma = G'_H(D)D_\sigma|_m / \{1 - G'_H(D)D_m\} \geq 0$*,* $x_\sigma|_h = \Xi_e W m_\sigma + \Xi_\sigma \geq 0$*, and* $Q_\sigma|_h = Q_x^M x_\sigma|_h + Q_\sigma^M|_x > 0$*.*

(b) The cash crowd-out. *The symmetric cash response is (29): cash rises in the conjecture if and only if precautionary loss relief exceeds marginal outside risk-premium incidence. With endogenous inventory, the authorized-settlement response is*

$$\widehat{Q}_\sigma = Q_x^M\{x_\sigma|_h + x_h h_\sigma^P\} + Q_\sigma^M|_x,$$
$$h_\sigma^P = W\lambda_\sigma^P, \qquad \lambda_\sigma^P = -\frac{\mathcal{F}_{H,\sigma}}{\mathcal{D}_H}, \tag{40}$$

*so* $\widehat{Q}_\sigma \geq 0$ *if and only if inventory accumulation does not absorb the combined price and routing response; a primitive sufficient condition is* $\mathcal{F}_{H,\sigma} \geq 0$*, and best-response stability alone does not sign* $\widehat{Q}_\sigma$*. At fixed* $(\sigma, W)$*, a tighter restriction changes the optimized issuer route only through its cash response:*

$$\widehat{Q}_\tau\Big|_{\sigma,W} = \widetilde{Q}_h h_\tau^P = \widetilde{Q}_h W \frac{\vartheta - \epsilon}{\mathcal{D}_H}. \tag{41}$$

*Because* $\widetilde{Q}_h \leq 0$*, the restriction lowers the route if and only if the marginal issuer-funded share exceeds payout ineligibility,* $\vartheta > \epsilon$*.*

(c) Claim migration and pass-through neutrality. *At fixed recovery and a smooth interior cash choice,*

$$w_\tau = \frac{\Delta_\tau^{W,P}}{\mathcal{D}_W}, \qquad \Delta_\tau^{W,P} = \epsilon + \lambda^P(\bar{\omega} - \epsilon) - \dot{\Delta}_{H,\lambda}\frac{\vartheta - \epsilon}{\mathcal{D}_H}. \tag{42}$$

*Claim migration is positive if and only if* $\Delta_\tau^{W,P} > 0$*. The allocation and this condition are independent of customer pass-through* $\varpi_H$ *under a flexible competitive fee.*

Part (b) is the procyclical-liquidity prediction previewed in the introduction. Part (c) disciplines a common intuition: because a competitive fee capitalizes pass-through one-for-one, whether the restriction pushes holders into wrappers turns on the eligibility loss $\epsilon$, the funded share of cash, and the funding-incidence term, and the advertised pass-through drops out. Two modeling choices behind these results are worth keeping distinct. An operator takes the aggregate price and conjecture as given yet internalizes how its own observable cash changes its customers' withdrawals; treating the sector as one price-impacting wrapper would mix competitive and strategic behavior. And first-loss incidence and foregone payout are transfers; counting them as social resource losses would double count them.

*Contracting robustness.* Pass-through neutrality is a long-run competitive-incidence result. Internet Appendix IA.II shows that a sticky fee restores customer pass-through to claim migration one-for-one, that a binding fee floor requires limited entry, and that an auditable covenant internalizing late-holder loss raises cash weakly while leaving issuer recovery governed by the same cash wedge as the baseline. These variants protect the mechanism.

# V. Two Balance Sheets, One Mechanism

At fixed inventory the stress mechanism is monotone: a worse conjecture lowers the value of ordinary-held tokens, raises the value of priority redemption, increases the tokens needed to fund par cash, and worsens late-claim dilution. Endogenous inventory may offset this effect, but outside investors simultaneously demand a larger risk rebate, which can crowd out precautionary cash and shift the funding of any given cash stock toward issuer redemption. This section builds the objects that aggregate these forces into a recovery map: the reduced-form contract, the cash wedge, and the two funding-flight terms.

Substituting claim clearing into the conditional equilibrium gives reduced cash, issuer scale, direct settlement, authorized settlement, and issuer value: writing $w := w(\sigma, \tau; \bar{Q})$, each reduced object below is a function of $(\sigma, \tau; \bar{Q})$,

$$h^r = h^P(\sigma, w, \tau), \qquad S^r = M - \mathcal{J}^{\text{fin}}(h^r, \sigma, \tau; \kappa_F), \qquad Z^r = Z^{\nu A}(\sigma, w, h^r), \tag{43}$$
$$Q^r = \widetilde{Q}(\sigma, w, h^r; \bar{Q}), \qquad V^r = V^I(w, h^r, \sigma, \tau; \kappa_F),$$

with smooth-point derivatives

$$h^r_\sigma = h^P_\sigma + h^P_W w_\sigma, \qquad h^r_\tau = h^P_\tau + h^P_W w_\tau, \qquad S^r_\sigma = -\varphi - \vartheta h^r_\sigma, \qquad S^r_\tau = -\eta_F - \vartheta h^r_\tau. \tag{44}$$

## A. Does Wrapper Liquidity Protect the Issuer?

Consider one more unit of wrapper cash at fixed $(\sigma, W, \tau)$. On the wrapper's balance sheet it is unambiguously protective. On the issuer's it does four things at once, and the issuer-implied haircut records their sum:

$$\mathcal{X}_h := \underbrace{\mathcal{H}_Q \widetilde{Q}_h}_{\substack{\text{route relief} \\ (\leq 0)}} + \underbrace{\mathcal{H}_V V_h^I}_{\substack{\text{franchise} \\ \text{cannibalization}}} \underbrace{- \vartheta\, \mathcal{H}_S}_{\substack{\text{scale: operating relief} \\ \text{minus capacity loss}}} + \underbrace{a_O(1-\vartheta)\mathcal{H}_Z}_{\substack{\text{perimeter} \\ \text{access load}}}. \tag{45}$$

Cash lowers routed authorized settlement in stress (route relief) and, by shrinking issuer scale through final debit, lowers scale-dependent operating commitments. Against these reliefs, its redemption-funded share sacrifices reserve margin; cash reduces franchise value if and only if (18) holds; its smaller scale removes feasible stress liquidity when issuer capacity binds (the scale term $-\vartheta\mathcal{H}_S$ nets operating relief, $-\vartheta\alpha\pi' q_R^* R_X'(S) \leq 0$, against capacity cannibalization, $-\vartheta\alpha\pi' q_S^* \geq 0$ when capacity binds), and its outside-placed share parks directly accessible inventory in the perimeter. *Cash stabilizes the issuer if and only if total relief exceeds total balance-sheet cannibalization*, compactly $\mathcal{X}_h < 0$.

When $\pi'(q) > 0$ and issuer liquidity capacity is slack, $q_S^* = 0$ and the condition reads $q_R^*\{\zeta_R(-\widetilde{Q}_h) + \vartheta R_X'(S)\} > q_V^* V_h^I + q_R^*\zeta_D a_O(1-\vartheta)$; when positive capacity binds, $q_V^* = 0$ and $q_S^* = -\overline{L}'(S)$, so $\zeta_R(-\widetilde{Q}_h) > \vartheta\{\overline{L}'(S) - R_X'(S)\} + \zeta_D a_O(1-\vartheta)$. The binding-capacity test has a transparent primitive form. With locally linear routing costs $a_B(q_B) = \beta_B q_B$, $c_R(Q) = \beta_R Q$, define $\theta_R := \beta_B/(\beta_B + \beta_R)$ and $\kappa_M := \beta_B\beta_R/(\beta_B + \beta_R)$; then $Q_x^M = \theta_R$, $\mathcal{B}_x = P - \kappa_M x$, and settlement relief is $\Omega_h := \zeta_R(-\widetilde{Q}_h) = \zeta_R\theta_R(1 - Wm_h)/(P - \kappa_M x)$. With the net scale mismatch $\delta_S(S) := \overline{L}'(S) - R_X'(S)$, cash stabilizes under binding positive issuer capacity exactly when

$$\Omega_h > \vartheta\delta_S(S) + \zeta_D a_O(1-\vartheta). \tag{46}$$

Settlement relief is stronger when more of a marginal sale is routed, net settlement is larger, marginal sale proceeds are lower, or cash more strongly reduces withdrawals; a steeper capacity schedule works against cash, while operating commitments that scale with the issuer work for it. The test depends on the *marginal* final-debit share $\vartheta$, not the average share $\bar{\omega}$, which governs levels. At $\vartheta = 1$ with no access the joint-balance-sheet benchmark obtains; at $\vartheta = 0$ issuer scale is locally fixed but extra cash still parks accessible inventory in the perimeter, so route relief must cover $\zeta_D a_O$.

A word on the sign of $\delta_S$. The mismatch is positive whenever the issuer's usable stress liquidity falls faster with scale than its operating commitments, as is natural for a reserve-backed issuer whose liquidity ladder is a portfolio pecking order: the marginal reserve dollar lost to pre-stress redemption is a liquid one (bank deposits, overnight repo, T-bills sold first), while operating commitments are dominated by fixed compliance and personnel costs. Then

$\overline{L}'(S)$ is near the liquid share of the marginal dollar and $R_X'(S)$ is near zero. The opposite sign is possible for an issuer whose costs scale strongly with size, and (46) prices that case too; the certified economy uses $\delta_S > 0$.

## B. Two Funding Flights and a Diversification Trade-Off

Now fix the cash stock and let beliefs or policy move only its *funding*. The recovery effect of one unit of final issuer debit, holding cash and claims fixed, is

$$\mathcal{X}_J := \xi(i - \mu_A)\mathcal{H}_V - \mathcal{H}_S - a_O\mathcal{H}_Z : \tag{47}$$

final debit shrinks issuer scale and removes perimeter inventory, including the direct settlement attached to that inventory. The belief- and policy-driven contributions to the recovery map are then

$$\Lambda_F := \mathcal{H}_V V_\sigma^{I,F} - \varphi\mathcal{H}_S - a_O\varphi\mathcal{H}_Z = \varphi\,\mathcal{X}_J, \qquad \Lambda_O := a_O'\mathcal{N}\mathcal{H}_Z \geq 0, \qquad \Pi_F := \eta_F\,\mathcal{X}_J. \tag{48}$$

Belief-driven flight $\Lambda_F$ multiplies the conjecture and therefore enters the recovery map's *gain*; policy flight $\Pi_F$ shifts the map itself: a tighter restriction makes payout-bearing outside investors require a larger placement adjustment and thereby migrates funding toward final debit at every conjecture. Access exercise adds $\Lambda_O$ to the belief loop but has no parallel policy term, because the access-cost distribution and fee are held fixed when $\tau$ changes. Both flight terms are distinct from $\mathcal{X}_h$: there the cash stock changes; here it is fixed while its financing migrates.

The quadratic family adds a closed-form depth result. When positive issuer capacity binds, $\mathcal{H}_V = 0$ and $-\mathcal{H}_S = \alpha\pi'\delta_S(S)$, so

$$\Lambda_F = \alpha\pi'\{\delta_S(S) - \zeta_D a_O\}\frac{\nu^2 p(1-a_O)\,\omega(1-\omega)}{\kappa_0}. \tag{49}$$

If $\delta_S(S) > \zeta_D a_O$, the loop is positive and *hump-shaped in relative depth*: it vanishes when either channel is locally unavailable, at full netting, or when access fully insures belief repricing, and is largest at $\omega = 1/2$. Funding diversification therefore lowers static reliance on either channel while maximizing the speed at which adverse beliefs migrate funding toward issuer redemption. We call this the flexibility–fragility trade-off; it lives entirely in the reallocation elasticity, which an exogenous funding share sets to zero.

PROPOSITION 2 (Liquidity substitution, finality, and two forms of funding flight): *Suppose Assumption 1 holds and consider a smooth regular point.*

(i) Static liquidity substitution. *An increase in wrapper cash lowers the issuer-implied haircut if and only if relief exceeds cannibalization, $\mathcal{X}_h < 0$. Under linear marginal routing costs and binding positive issuer capacity, this condition is exactly (46). The static balance-sheet cost of cash is governed by the effective marginal final-debit share $\vartheta$; directly accessible perimeter inventory adds a separate settlement load; the average share $\bar{\omega}$ governs levels, not this local test.*

(ii) Belief- and policy-induced funding flight. *At an interior funding split, the reallocation elasticities of (9) are $\mathcal{J}_\sigma^{\text{fin}} = \varphi = \nu^2 p(1 - a_O)/\mathcal{D}_F \geq 0$ and $\mathcal{J}_\tau^{\text{fin}} = \eta_F = \nu^2/\mathcal{D}_F \geq 0$, writing $\mathcal{D}_F := (C_I^0)'' + (C_O^0)''$ for funding curvature; belief-driven reallocation is strict when $\nu > 0$ and $a_O < 1$, policy-driven reallocation when $\nu > 0$. The fixed-cash recovery effects are $\Lambda_F$ and $\Pi_F$ in (48). When positive issuer capacity binds,*

$$\Lambda_F = \alpha\pi'\varphi\{\delta_S(S) - \zeta_D a_O\}, \qquad \Pi_F = \alpha\pi'\eta_F\{\delta_S(S) - \zeta_D a_O\}. \tag{50}$$

*Both worse beliefs and a tighter restriction weaken recovery through reallocation exactly when the capacity-minus-operating-load slope exceeds the direct-settlement load on perimeter inventory; beliefs additionally induce the exercise loop $\Lambda_O$.*

(iii) Three statistics and two shock loadings. *The average share $\bar{\omega}$ governs balance-sheet levels, $\vartheta$ governs static cash-induced scale loss, and $\varphi$ governs the belief-driven funding slope; policy reaches the same reallocation margin without access insurance, with loading $\eta_F = \varphi/[p(1-a_O)]$ when $a_O < 1$. In the interior quadratic family (15), $\vartheta = \nu\omega$, $\varphi = \nu^2 p(1 - a_O)\,\omega(1 - \omega)/\kappa_0$, and $\eta_F = \nu^2\omega(1 - \omega)/\kappa_0$. Both reallocation effects vanish at persistent single-channel corners and at full netting, and both are largest at balanced relative depth. Full effective access eliminates belief-induced reallocation but not policy-driven reallocation or direct settlement from perimeter inventory.*

This is the paper's liquidity-substitution result: more cash can protect the wrapper yet weaken issuer recovery, and a yield restriction can simultaneously change how much cash the wrapper holds and how that cash is funded, with independent signs. The next two sections characterize how the belief slope is amplified into equilibrium. Figure 3 previews the object they study: the recovery map whose gain the funding market controls and whose position the yield restriction shifts.

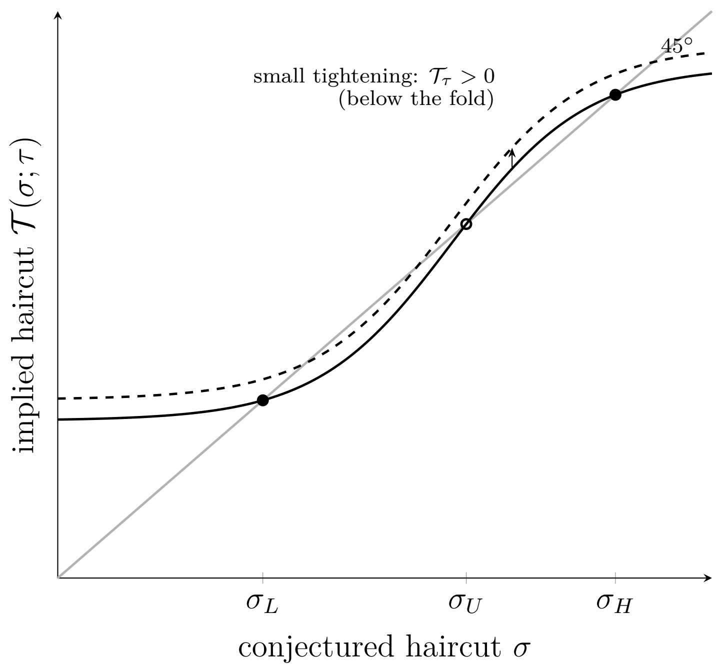


Figure 3: The recovery map (schematic; positions of the three crossings match the certified economy of Section VII qualitatively). Filled circles are the stable safe and fragile equilibria $\sigma_L$ and $\sigma_H$; the open circle is the unstable middle equilibrium $\sigma_U$. The map's slope at a crossing is the loop gain $\mathcal{T}_\sigma$, which the funding market raises through belief-driven flight $\Lambda_F$; a tighter yield restriction shifts the map by the forcing $\mathcal{T}_\tau$, whose funding component is $\Pi_F$. The dashed curve shows a small tightening that remains below the first fold: all three crossings survive, the safe and fragile equilibria move up, and the unstable separator moves down, shrinking the safe basin.

# VI. A Unit-Free Phase Diagram for Funding Fragility

This section isolates the feedback loop the funding market adds and shows that its global structure is governed by two dimensionless numbers. The mechanism is direct. More issuer-facing funding completes more settlement, which raises recovery pressure; a worse recovery raises the outside quote, which tilts funding further toward the issuer. Whether this loop merely amplifies or creates multiple self-fulfilling configurations depends on a race between the loop's local gain and the congestion that builds as the funding mix becomes lopsided. Two statistics summarize the race: *feedback leverage*, the loop gain at the balanced mix, and *route tilt*, the signed profitability advantage of issuer-facing funding measured in congestion units.

## A. An Exact Benchmark

To isolate the loop, consider an embedded benchmark in which an audited liquidity covenant fixes $h > 0$, claim scale, and finality $\nu \in [0, 1]$, while the split between issuer redemption and outside placement remains endogenous. Use the quadratic funding family, maintain an interior split, suppose positive issuer capacity binds with constant net scale sensitivity $\delta_S > 0$, and let the resolution technology be affine on the reached range. Recall $d(\omega)$ from (14) and define $\beta_S := \alpha\pi'\delta_S > 0$. All recovery forces other than funding reallocation are summarized by an affine baseline with loop gain $\Lambda_0 \in [0, 1)$ and direct policy relief $\upsilon \geq 0$. The exact benchmark map is

$$\mathcal{T}^B_{\omega,\nu}(\sigma; \tau) := \operatorname{proj}_{[0,\bar{\sigma}]}\{\sigma_B + \Lambda_0\sigma - \upsilon\tau + \beta_S\left[\nu\omega h + \nu d(\omega)\{\nu[p\sigma - i(\tau)] - \zeta_N(\nu)\}\right]\}\,, \tag{51}$$

where proj maps an implied haircut onto its feasible interval. This is an exact subclass of the full model: each extra unit of final date-0 redemption raises the implied haircut by $\beta_S$, and variation in $\omega$ re-solves both the funding allocation and the map, so the flight term is an equilibrium object throughout. The factor $\nu^2$ is the double-finality loading of (9). Settlement-credit cost shifts the level of tender and hence the map's intercept, while the loop gain, which depends on how strongly the quote responds to beliefs, is unchanged by a conjecture-independent per-tender charge. For later use, write the unprojected intercept and gain as

$$A_{\omega,\nu}(\tau) := \sigma_B - \upsilon\tau + \beta_S\{\nu\omega h - \nu^2 d(\omega) i(\tau) - \nu d(\omega)\zeta_N(\nu)\}, \qquad \Lambda^B_{\omega,\nu} := \Lambda_0 + \beta_S p\nu^2 d(\omega), \tag{52}$$

so $\mathcal{T}^B_{\omega,\nu} = \operatorname{proj}_{[0,\bar{\sigma}]}\{A_{\omega,\nu}(\tau) + \Lambda^B_{\omega,\nu}\sigma\}$.

## B. Feedback Leverage and Route Tilt

Hold $(h, \omega, \tau, d)$ fixed, let $\mathcal{J}_0 = \omega h$, and retain an interior allocation with $\rho_F \geq 0$. To let the balance-sheet link be nonlinear, let the reached range of a three-times continuously differentiable monotone recovery backbone contain the right side of

$$\mathcal{R}_B(\sigma^B; \tau) = \beta_S\nu\mathcal{J}, \qquad \mathcal{R}_{B,\sigma}(\sigma^B; \tau) > 0. \tag{53}$$

For $j \in (0, h)$, let $\sigma(j, \nu)$ solve $\mathcal{R}_B(\sigma(j, \nu); \tau) = \beta_S\nu j$. At the balanced mix, write $\sigma_0(\nu) := \sigma(\mathcal{J}_0, \nu)$ and define the settlement-adjusted funding tilt $\Delta^G_N(\nu) := \nu\{p\sigma_0(\nu) - i(\tau)\} - \zeta_N(\nu)$.

At an equilibrium allocation $\mathcal{J}(\nu,d)$, write $\sigma^B(\nu,d) := \sigma(\mathcal{J}(\nu,d),\nu)$ and define the associated outside quote by $\rho_F(\nu,d) := p\sigma^B(\nu,d) - i(\tau)$. The affine benchmark has $\mathcal{R}_B = B_0\sigma - (\sigma_B - \upsilon\tau)$ with $B_0 := 1 - \Lambda_0$, and hence $\Delta_N^A(\nu) := B_0\Delta_N^G(\nu) = \nu\{p[\sigma_B - \upsilon\tau + \beta_S\nu\mathcal{J}_0] - B_0 i(\tau)\} - B_0\zeta_N(\nu)$. For the convex proportional-depth technology (13), define the reduced architecture residual

$$\mathcal{G}_F(j,\nu;d) := \frac{\Phi_{F,j}(j;h,\omega)}{d} - \nu\{p\sigma(j,\nu) - i(\tau)\} + \zeta_N(\nu), \tag{54}$$

and write $r_N(\nu,d) := \nu\rho_F(\nu,d) - \zeta_N(\nu) = \Phi_{F,j}(\mathcal{J};h,\omega)/d$ for the funding tilt at the allocation. Every interior equilibrium has $\mathcal{G}_F = 0$, and, with $\Phi_{F,jj}$ and $\mathcal{R}_{B,\sigma}$ evaluated at the equilibrium allocation, the net curvature is

$$\mathcal{G}_{F,j} = \mathcal{M}_F := \frac{\Phi_{F,jj}}{d} - \frac{\beta_S p\nu^2}{\mathcal{R}_{B,\sigma}} = \frac{\Phi_{F,jj}}{d}\{1 - \Lambda_F^G\}, \qquad \Lambda_F^G := \frac{\beta_S p\nu^2 d}{\mathcal{R}_{B,\sigma}\Phi_{F,jj}}, \tag{55}$$

so positive net curvature is exactly the local recovery-stability margin under both direct iteration and convex revision: $\mathcal{M}_F > 0$.

The simplest globally stabilizing route technology is locally quadratic but more congested in the tails. For $a_F, b_F, \bar{\Phi}_F > 0$ and $x = j - \mathcal{J}_0$, let

$$\Phi_F(j) = \bar{\Phi}_F + \frac{a_F}{2}x^2 + \frac{b_F}{12}x^4, \qquad \mathcal{R}_B = B_0\sigma - (\sigma_B - \upsilon\tau). \tag{56}$$

Define the dimensional feedback–congestion gap, turning-point distance, and congestion buffer,

$$k_F(\nu) := \frac{\beta_S p\nu^2}{B_0} - \frac{a_F}{d}, \qquad s_F(\nu) := \sqrt{\frac{d\,k_F(\nu)}{b_F}}, \qquad \overline{\Delta}_F(\nu) := \frac{2}{3}k_F(\nu)^{3/2}\sqrt{\frac{d}{b_F}}, \tag{57}$$

and, for normalized funding imbalance $y_F := \sqrt{b_F/a_F}\,x$, the two dimensionless indices and cubic normal form

$$\mathrm{FL}(\nu) := \frac{\beta_S p\nu^2/B_0}{a_F/d}, \qquad \mathrm{RT}(\nu) := \frac{3\sqrt{b_F/d}}{2(a_F/d)^{3/2}}\Delta_N^G(\nu), \tag{58}$$

$$\frac{3\sqrt{b_F/d}}{(a_F/d)^{3/2}}\mathcal{G}_F(j,\nu;d) = y_F^3 - 3\{\mathrm{FL}(\nu) - 1\}y_F - 2\,\mathrm{RT}(\nu). \tag{59}$$

Feedback leverage is exactly the recovery-map gain at the balanced mix, while route tilt is the signed funding advantage in curvature units. The positive normalization leaves roots unchanged, so, subject to endpoint and interiority conditions, the benchmark phase depends only on $(\mathrm{FL}, \mathrm{RT})$, not on measurement units. When $\mathrm{FL} \leq 1$, the residual is monotone; above one, tail congestion restores stability away from the balanced mix, and the normalized fold semi-width is $(\mathrm{FL}-1)^{3/2}$. The dimensional objects $s_F$ and $\overline{\Delta}_F$ locate those folds inside the economic domain.

Theorem 1 (Final settlement and funding fragility): *Fix a compact finality interval $I_\nu = [\underline{\nu}, \overline{\nu}] \subset (0,1)$.*

(i) Primitive congestion benchmark. *Under (56), suppose, for every $\nu \in I_\nu$, that*

$$\mathrm{FL}(\nu) > 1, \qquad s_F(\nu) < \min\{\mathcal{J}_0, h - \mathcal{J}_0\}, \qquad \mathcal{G}_F(0,\nu) < 0 < \mathcal{G}_F(h,\nu). \tag{60}$$

*There are exactly three interior simple equilibria if and only if*

$$|\mathrm{RT}(\nu)| < \{\mathrm{FL}(\nu) - 1\}^{3/2} \quad \Longleftrightarrow \quad |\Delta_N^G(\nu)| < \overline{\Delta}_F(\nu). \tag{61}$$

*The unit-free multiplicity band has width $2(\mathrm{FL}-1)^{3/2}$ and elasticity $3\mathrm{FL}/[2(\mathrm{FL}-1)]$ with respect to feedback leverage. The outer equilibria are locally stable and the middle equilibrium is unstable. In addition, suppose*

$$\mathrm{RT}'(\nu) > \left|\frac{\mathrm{d}}{\mathrm{d}\nu}\{\mathrm{FL}(\nu)-1\}^{3/2}\right|, \quad \mathrm{RT}(\underline{\nu}) < -\{\mathrm{FL}(\underline{\nu})-1\}^{3/2}, \quad \mathrm{RT}(\overline{\nu}) > \{\mathrm{FL}(\overline{\nu})-1\}^{3/2}. \tag{62}$$

*In the affine benchmark, the left derivative is proportional to $p\sigma_0 - i + m_N + p\beta_S \nu \mathcal{J}_0/B_0$, while the band derivative is $3\mathrm{FL}\sqrt{\mathrm{FL}-1}/\nu$: settlement-value and cost-saving drift must dominate band expansion. Then there are unique thresholds*

$$\mathrm{RT}(\nu_\downarrow) = -\{\mathrm{FL}(\nu_\downarrow) - 1\}^{3/2}, \qquad \mathrm{RT}(\nu_\uparrow) = \{\mathrm{FL}(\nu_\uparrow) - 1\}^{3/2}, \qquad \underline{\nu} < \nu_\downarrow < \nu_\uparrow < \overline{\nu}. \tag{63}$$

*Below $\nu_\downarrow$ the unique equilibrium is on the low-tender branch; between the thresholds, stable low- and high-tender equilibria coexist with an unstable middle equilibrium; above $\nu_\uparrow$ the unique equilibrium is on the high-tender branch. The high and middle branches meet at $\nu_\downarrow$, the low and middle branches at $\nu_\uparrow$; both contacts are generic folds.*

(ii) Convex revision, selection, and hysteresis. *Let the increasing conditional recovery map $\mathcal{T}_F$ preserve a compact interval. Between quote rounds the facility's risk desk revises the posted recovery haircut $\widehat{\sigma}_n$ used in placement and collateral quotes, trading a convex cost of moving* *the posted quote and margin schedule against a convex pricing and capital loss from missing the recovery target implied by current funding; Internet Appendix* (IA.177)–(IA.180) *state the protocol class. Every protocol in the class moves the posted haircut toward its target without overshooting and by at least a fixed fraction of the gap, and it has exactly the fixed points of* $\mathcal{T}_F$*. Inside the multiplicity region it therefore selects by basin: posted haircuts below the middle root converge to* $\sigma_L$*, those above it to* $\sigma_H$*, and quadratic costs give partial adjustment as a special case. Under quasistatic warm starts, increasing finality preserves the low branch until* $\nu_\uparrow$ *and decreasing finality preserves the high branch until* $\nu_\downarrow$*:* $(\nu_\downarrow, \nu_\uparrow)$ *is the exact hysteresis region for every persistent facility-repricing protocol in this class.*

(iii) Local fold geometry. *The global conditions above are sufficient, not necessary. Suppose an interior pair* $(\mathcal{J}_f, \nu_f) \in (0, h) \times (0, 1)$ *satisfies*

$$\mathcal{G}_F = \mathcal{G}_{F,j} = 0, \qquad C_f := \mathcal{G}_{F,jj} \neq 0, \qquad D_f := -\mathcal{G}_{F,\nu} > 0. \tag{64}$$

*Then* $(\mathcal{J}_f, \nu_f)$ *is a generic fold: on the side* $C_f(\nu - \nu_f) > 0$*, two local branches separate at the square-root rate; the branch with* $C_f(\mathcal{J} - \mathcal{J}_f) > 0$ *is stable and the other unstable; any distinct hyperbolic equilibrium persists. Settlement cost postpones a fold, stronger recovery exposure brings it forward, and depth brings it forward when funding tilts toward issuer-facing tender.*

Internet Appendix Theorem IA.1 extends the identical one–three–one phase diagram, root stability, and fold ordering to the full single-valley net-curvature class, including asymmetric route costs and nonlinear recovery backbones, under endpoint and single-valley conditions on $(\mathcal{G}_F, \mathcal{M}_F)$ alone. Conditional issuer and funding choices remain unique throughout that class; system feedback alone makes the reduced residual nonmonotone. Selection in part (ii) is robust within the convex-revision class, and it carries an institutional price: the regulator cannot choose a root costlessly.

Figure 4 separates the unit-free phase restriction from one finality path through it.

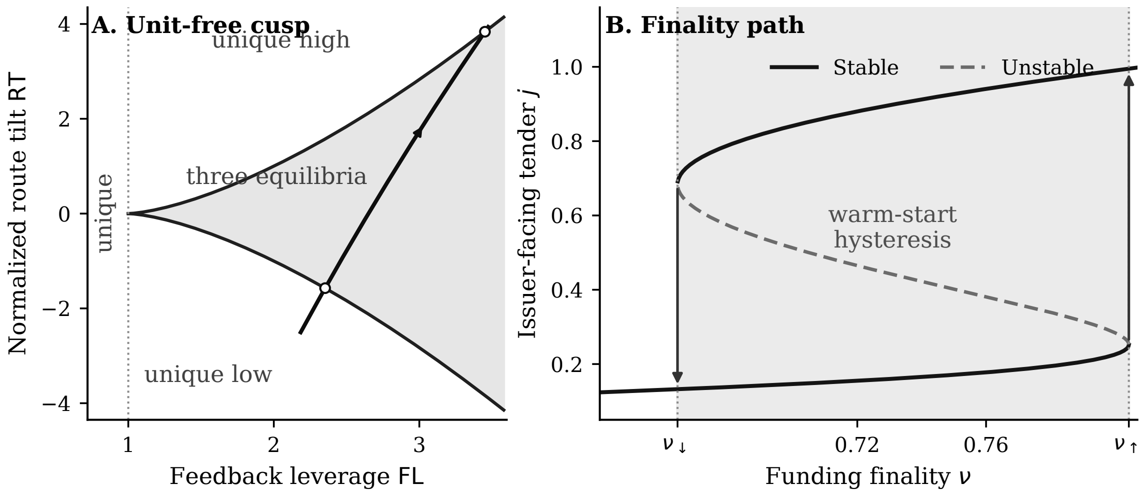


Figure 4: Unit-free phase diagram. Panel A shades the three-equilibrium cusp $|\mathrm{RT}| < (\mathrm{FL}-1)^{3/2}$; the black path is the certified economy as finality rises, with circles at its folds. Panel B shows its stable (solid) and unstable (dashed) tender branches and warm-start jumps. Section VII.B supplies magnitudes.

## C. Observable Content

The theorem maps its two statistics to route-level data. Let $\mathcal{J}^q(\widehat{\sigma})$ denote the unique conditional issuer-facing allocation at an interior funding choice. In the congestion benchmark,

$$\frac{\partial \mathcal{J}^q}{\partial \widehat{\sigma}} = \frac{d\,\nu p}{a_F + b_F(\mathcal{J}^q - \mathcal{J}_0)^2} > 0, \qquad \Lambda_F^G = \frac{\beta_S p \nu^2 d}{B_0\{a_F + b_F(\mathcal{J} - \mathcal{J}_0)^2\}}, \tag{65}$$

$$\mathrm{FL}(\nu) = \Lambda_F^G\Big|_{\mathcal{J}=\mathcal{J}_0} = \frac{\beta_S \nu}{B_0}\,\frac{\partial \mathcal{J}^q}{\partial \widehat{\sigma}}\bigg|_{\mathcal{J}^q=\mathcal{J}_0}. \tag{66}$$

Given exogenous recovery-news variation and matched tender-to-debit records, feedback leverage can be estimated from the issuer-facing route response and recovery exposure. Estimating route tilt additionally requires quotes on both sides of the balanced funding mix, local support, and the maintained quartic-congestion specification. A scalar phase diagnostic is the normalized cusp slack

$$\mathrm{CS}(\nu) := \{\mathrm{FL}(\nu) - 1\}^3 - \mathrm{RT}(\nu)^2, \tag{67}$$

positive exactly on the three-equilibrium region, negative on the unique-equilibrium region, and zero on a fold boundary or the cusp tip. Away from the tip, the gradient $(3(\mathrm{FL}-1)^2, -2\mathrm{RT})$

maps an estimated covariance into a delta-method standard error; at the tip it vanishes and a joint confidence set is required.

Four implications follow. First, adverse recovery news shifts a fixed cash requirement toward issuer-facing funding, most strongly near a balanced mix; the cusp is a cross-market mechanism restriction that aggregate issuance or redemption series cannot test. Second, completed issuer debit, rather than gross tender or chain-wide burns, is the treatment-relevant settlement quantity: exposure scales with $\nu$ but the reallocation loop scales with $\nu^2$, so valid measurement must match each gross tender to the amount finally debited. Third, testing hysteresis requires repeated changes in the same architecture, with separate upward and downward event time; a single policy event or pooled directions cannot reveal the fold path. Fourth, claim migration is neither necessary nor sufficient for a recovery effect (equation (82) below), so a mechanism test must observe wrapper balances and cash, accepted redemptions, perimeter inventory, route prices, and final debits jointly. The filing exhibit in the Internet Appendix establishes the scale of issuer–platform transfers, and Section VII.B reports the public counterparts of these objects for the March 2023 window.

# VII. The Full Contract: Multiplicity and the Policy Reversal

The benchmark held cash and claims fixed by covenant. We now let every margin adjust. The full recovery map composes the haircut map with the entire reduced-form contract,

$$\mathcal{T}(\sigma;\tau,\bar{Q}) := \mathcal{H}(Z^r, Q^r, V^r, S^r), \tag{68}$$

and two objects organize the analysis: its slope and its shift. At a smooth point, the total loop gain is

$$\begin{aligned}\mathcal{T}_\sigma = {} & \mathcal{H}_Z(\widehat{Z}_\sigma + \widehat{Z}_W w_\sigma) \\ & + \mathcal{H}_Q(\widetilde{Q}_\sigma + \widetilde{Q}_W w_\sigma) + \mathcal{H}_V V_W^I w_\sigma + \Lambda_F + \Lambda_O + \mathcal{X}_h\, h_\sigma^r.\end{aligned} \tag{69}$$

The first line is the original-holder loop; the second contains the fixed-cash wrapper route, recovery-induced claim exit, the funding-flight loop, the perimeter access-exercise loop, and the joint-balance-sheet effect of endogenous cash, which combines precautionary relief with the crowd-out of (29). Holding the conjecture fixed, a tighter restriction shifts the map by

$$\mathcal{T}_\tau = \left(\mathcal{H}_Z \widehat{Z}_W + \mathcal{H}_Q \widetilde{Q}_W + \mathcal{H}_V V_W^I\right) w_\tau + \mathcal{H}_V V_\tau^{I,0} + \Pi_F + \mathcal{X}_h\, h_\tau^r. \tag{70}$$

Defining the fixed-cash recovery exposure of claim scale

$$\mathcal{X}_W := \mathcal{H}_Z \widehat{Z}_W + \mathcal{H}_Q \widetilde{Q}_W + \mathcal{H}_V V_W^I, \tag{71}$$

equations (33) and (44) give the primitive intensive–extensive decomposition

$$\mathcal{T}_\tau = \left(\mathcal{X}_W + \mathcal{X}_h h_W^P\right) \frac{\Delta_\tau^{W,P}}{\mathcal{D}_W} + \mathcal{H}_V V_\tau^{I,0} + \Pi_F + \mathcal{X}_h W \frac{\vartheta - \epsilon}{\mathcal{D}_H}. \tag{72}$$

Each of the four forces matches a different regulator's intuition: claim migration (including cash carried by new claims), direct payout-cost relief, policy-induced funding flight, and the cash response on outstanding claims. The theorem below determines which combinations survive equilibrium.

Theorem 2 (Settlement architecture, recovery amplification, and yield policy): *For the general primitive model in parts* (i) *and* (iii)*, suppose Assumption 1 holds; the benchmark and constructive family in part* (ii) *use the conditions stated there and are verified directly from primitives.*

(i) Recovery equilibrium, stability, and selection. *For every $(\tau, \bar{Q})$, the recovery map has a fixed point; at every smooth point its slope is (69). Let $\Lambda_{-FO} := \mathcal{T}_\sigma - \Lambda_F - \Lambda_O$. At a point where positive issuer capacity binds and the map is locally increasing, if $|\Lambda_{-FO}| < 1$, funding reallocation plus access exercise overturns local stability exactly when*

$$\frac{\nu^2 p(1 - a_O)}{\mathcal{D}_F}\{\delta_S(S) - \zeta_D a_O\} + \zeta_D a_O' \mathcal{N} > \frac{1 - \Lambda_{-FO}}{\alpha \pi'}. \tag{73}$$

*In the full-finality, no-access benchmark $(\nu, a_O, a_O') = (1, 0, 0)$ this reduces, when $\delta_S(S) > 0$, to*

$$\varphi > \frac{1 - \Lambda_{-FO}}{\alpha \pi' \delta_S(S)}, \tag{74}$$

*equivalently, in the interior quadratic family, $\kappa_0 < p\omega(1-\omega)\alpha\pi'\delta_S(S)/(1-\Lambda_{-FO})$ holding the remaining local objects fixed. These are local sufficient-statistic comparisons; a gain above one somewhere is necessary but not sufficient for multiplicity, since the residual must also cross* *zero in the required order. A weakly increasing map has a nonempty complete lattice of fixed points reached by endpoint iteration; a uniform bound* $|\mathcal{T}_\sigma| \le \kappa_R < 1$ *gives a unique globally attracting equilibrium; a differentiable root is locally stable when* $|\mathcal{T}_\sigma| < 1$ *and unstable when* $|\mathcal{T}_\sigma| > 1$*. With an increasing map, exactly three simple roots, and residual signs* $+, -, +, -$*, the outer roots are stable and the middle root unstable, and every persistent convex-revision protocol in Theorem 1*(ii) *selects the low root from estimates below the middle root and the high root from estimates above it, with quasistatic warm starts jumping only when the selected branch disappears. The same conclusions hold on a closed invariant subinterval, including* $[0, 0.20]$ *in the constructive family.*

(ii) Netting, multiplicity, and the stable reversal. *In the benchmark of Section VI, fix an interior* $\omega$ *and* $\nu \in [0,1]$*. An interior mixed allocation strictly lowers private funding cost relative to forcing all cash through either route at the same channel depths, and its recovery loop gain is exactly* $\Lambda^B_{\omega,\nu}$ *in* (52)*; settlement-credit parameters shift* $A_{\omega,\nu}$ *but not this gain. If* $\Lambda^B_{\omega,\nu} < 1$*, the projected map is a contraction with a unique globally stable fixed point. If* $\Lambda^B_{\omega,\nu} > 1$ *and*

$$(1 - \Lambda^B_{\omega,\nu})\bar{\sigma} < A_{\omega,\nu}(\tau) < 0, \tag{75}$$

*the fixed-point set is exactly*

$$\left\{0, -\frac{A_{\omega,\nu}(\tau)}{\Lambda^B_{\omega,\nu} - 1}, \bar{\sigma}\right\}, \tag{76}$$

*with stable boundary equilibria and an unstable interior equilibrium.*

*Both single-channel limits have loop gain* $\Lambda_0 < 1$*. Define the critical funding finality*

$$\nu_c := 2\sqrt{\frac{\kappa_0(1-\Lambda_0)}{\beta_S p}}. \tag{77}$$

*For* $\nu < \nu_c$*,* every *relative-depth design is a contraction; if* $\nu_c > 1$*, every feasible finality–depth design is a contraction. If* $\nu > \nu_c$*, define*

$$\omega^S_\pm(\nu) := \frac{1}{2}\left(1 \pm \sqrt{1 - \frac{4\kappa_0(1-\Lambda_0)}{\beta_S p \nu^2}}\right): \tag{78}$$

*whenever an interior fixed point exists, it is unstable exactly for* $\omega \in (\omega^S_-(\nu), \omega^S_+(\nu))$*, the map is a contraction outside* $[\omega^S_-, \omega^S_+]$*, and the instability interval expands strictly with* $\nu$*. Within this benchmark, access to both routes can lower cost while creating safe–fragile multiplicity absent from either single channel, and netting eliminates that architecture-induced fragility* *without closing either channel. At a stable interior benchmark fixed point,*

$$\sigma_\tau^B = \frac{\beta_S \nu^2 d(\omega) - \upsilon}{1 - \Lambda_0 - \beta_S p \nu^2 d(\omega)} : \tag{79}$$

*a tighter restriction worsens recovery if and only if funding reallocation overturns its direct benefit,* $\beta_S \nu^2 d(\omega) > \upsilon$*. The set of relative depths satisfying both the strict stability and reversal inequalities is nonempty if and only if*

$$\upsilon < \min\left\{\frac{\beta_S \nu^2}{4\kappa_0}, \frac{1-\Lambda_0}{p}\right\}, \tag{80}$$

*and any such depth exhibits a stable interior reversal whenever its intercept places the fixed point strictly inside* $(0, \bar{\sigma})$*.*

*The certified open set shows that the policy reversal can survive full contract endogeneity. The full contract can also sustain three roots, with the certified multiplicity driven by the routing-and-claims loop. The threshold* $\nu_c$ *is a benchmark result; in the full contract, uniqueness and multiplicity are governed by part* (i)*. Internet Appendix Section IA.VI certifies a relatively open primitive family in which the full nonlinear contract has exactly three genuine competitive equilibria on unrestricted action domains, with strictly positive cash and claim scale, both funding routes active, and a finality margin strictly inside* $(0, q_A f_A)$*; at every safe root, funding reallocation reverses the sign of the remaining policy forces. The same appendix certifies two positive-area* $(\nu, a_O)$ *rectangles with direct perimeter access and the same safe–unstable–fragile ordering: the high-finality rectangle exhibits policy reversal, whereas lower finality restores the intended policy sign without removing multiplicity or access. The exact sets, strict margins, and outward-rounded global-optimality certificates are in* (IA.58)–(IA.61)*; the certified roots exhaust the unrestricted contract. The construction establishes existence on an open set of primitives.*

(iii) General yield-policy effect. *At a smooth stable fixed point,*

$$\sigma_\tau = \frac{\mathcal{T}_\tau}{1 - \mathcal{T}_\sigma}, \tag{81}$$

*with numerator exactly* (72)*: claim migration has the sign of* $\Delta_\tau^{W,P}$ *and cash the sign of* $\vartheta - \epsilon$*; neither is unconditionally positive. At a point with zero net migration, the uncollected forcing is*

$$\mathcal{T}_\tau = \mathcal{H}_V V_\tau^{I,0} + \Pi_F + \mathcal{X}_h W \lambda_\tau^P : \tag{82}$$

*yield policy can change recovery without any claim migration, through direct payout cost, the funding composition of cash, and cash intensity. If the map is increasing throughout a policy rectangle and $\mathcal{T}_\tau$ has one weak sign throughout, both the least and greatest equilibria move in that direction; a sign at one selected root does not order the correspondence globally.*

The theorem links the two balance sheets through the funding statistics. Familiar proxies map poorly into the mechanism: claim migration can rise while recovery improves, wrapper cash can weaken the issuer it appears to protect, and a diversified funding mix maximizes the very reallocation elasticity that transmits bad news. What governs outcomes are the marginal funding share, for the static cash effect, and the flight elasticity, for dynamic amplification. In the embedded design benchmark, finality determines whether the funding-instability region exists and how wide it is. The netting threshold $\nu_c$ makes the determinants explicit: it rises with funding congestion $\kappa_0$ and with the baseline loop's distance from criticality, and falls with recovery exposure $\beta_S$ and stress probability $p$, so deep, cheap funding markets attached to fragile issuers require the most netting.

## A. Separating the Two Loops

The certified economy separates the mechanisms in (69)–(72). Figure 5 reports the outward-rounded midpoint decompositions at the three certified roots $\sigma = (0.0713, 0.1559, 0.1839)$.

The *multiplicity* is created by the classic routing-and-claims loop. At the safe root the total gain is 0.118, with every component small; at the middle root the gain is 1.225, of which 1.182 is the fixed-cash routing and claim-exit loop (withdrawals, forced sales, and priority routing), while funding flight contributes only 0.028 and endogenous cash $-0.030$.

The *policy reversal* belongs to the funding loop. At the safe root the forcing decomposes into claim-route 0.000, funding flight $\Pi_F = 0.081$, and endogenous cash $-0.007$: absent the funding channel the restriction would (weakly) improve recovery at the safe root, and funding flight alone flips the sign, delivering $\sigma_\tau = \mathcal{T}_\tau/(1 - \mathcal{T}_\sigma) = 0.084 > 0$. At the fragile root the forcing is 0.164 with amplification $1/(1 - 0.775)$, giving $\sigma_\tau = 0.732$: the same restriction is almost nine times as damaging where feedback is strong. (At the unstable middle root the formal expression is $-1.118$, which is not a policy object.)

The decomposition suggests distinct empirical tests: a claim-only model can match the number of equilibria yet miss the sign of yield policy, and a funding-only model has the opposite blind spot.

*Computational certification.* All reported roots, margins, and decompositions are verified

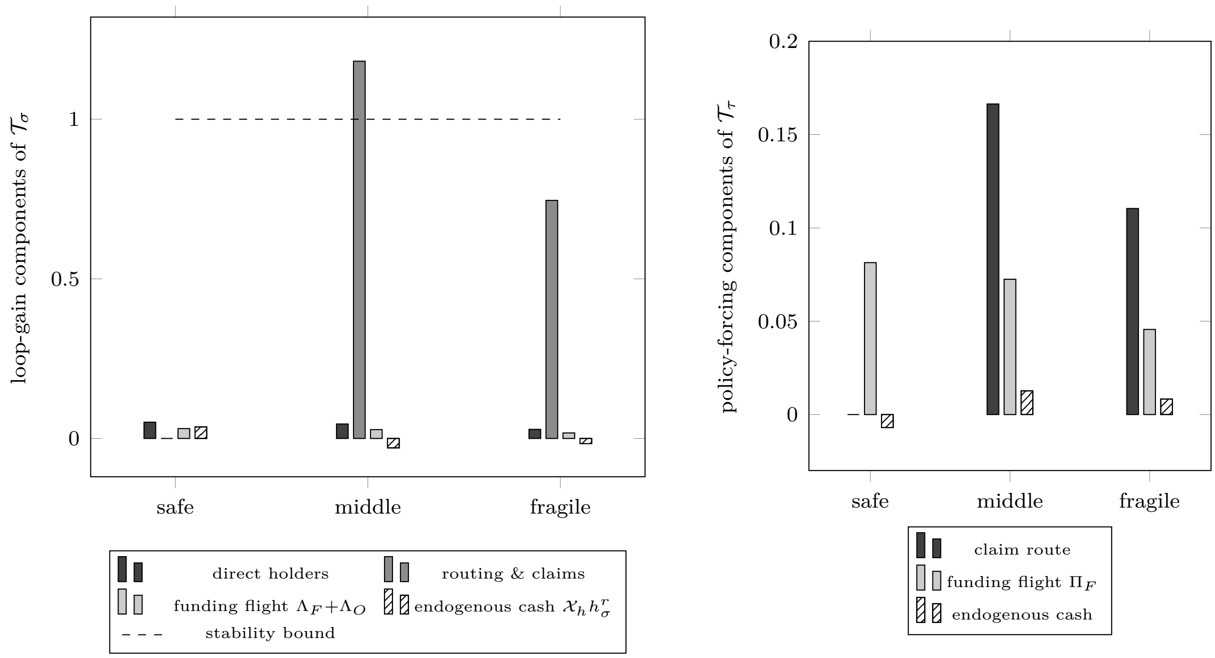


Figure 5: Anatomy of the certified economy at its three roots (outward-rounded midpoints; Internet Appendix IA.VI). Left: components of the loop gain $\mathcal{T}_\sigma$; totals 0.118, 1.225, 0.775. The middle root's instability is driven by the routing-and-claims loop. Right: components of the policy forcing $\mathcal{T}_\tau$; totals 0.074, 0.252, 0.164. At the safe root the claim route contributes zero and endogenous cash is stabilizing, so funding flight $\Pi_F$ alone makes the restriction harmful, $\sigma_\tau = 0.084 > 0$.

with outward-rounded interval arithmetic (Krawczyk existence–uniqueness operators on 256-bit ball arithmetic): each root lives in a rigorously enclosed box; a sign-certified exclusion sweep covers the entire unrestricted action space, so no fourth equilibrium exists; and every strict inequality in Theorem 2(ii) holds with certified margin on a relatively open primitive neighborhood. Multiplicity and the reversal therefore hold with certified margins, beyond the reach of grid or tolerance artifacts. Internet Appendix IA.VI contains the construction, the enclosures, and the reproducibility protocol.

## B. How Large Is the Funding Channel? A Calibrated Illustration

The certified economy establishes existence. This subsection measures the funding channel's inputs at the scale of the largest regulated issuer. Three public data layers enter: minute-level USDC prices and daily on-chain burns around March 2023; the monthly portfolio of the money market fund that holds the noncash portion of USDC reserves; and platform terms from company filings and archived reward pages. Table I collects the inputs. The exercise fixes levels and event moments from public sources under the maintained functional forms, and it reports ranges where public data end.

*Scale and payout.* The wrapper layer is first order. Against \$75.27 billion of USDC at end-2025, mapped on-chain wrappers held \$13.31 billion at the October 2025 cross-section and the affiliated platform up to \$12.50 billion more, so $W/M$ lies between roughly 0.18 and 0.34. On the payout side, the reserve fund's seven-day gross yield averaged 4.95, 5.18, 4.26, and 3.69 percent over 2023 to 2026, and the platform's advertised top-tier USDC reward tracked it nearly one for one, financed by an agreement that transferred 53.1 percent of reserve income in 2025. A top-tier pass-through near one puts the economy in exactly the regime of Proposition 1(c): migration runs through eligibility and funding incidence, with the advertised rate capitalized into fees. In policy units, the pre-GENIUS regime is $\tau \approx 0$, and a full prohibition moves $\tau$ by 370 to 520 basis points depending on the year's rate level.

*The reserve ladder and $\delta_S$.* Across 42 monthly filings, the marginal reserve dollar is held at the extreme short end: over 2025 to 2026, 67.2 percent of fund holdings matured within seven days (63.0 percent in overnight-dominated repurchase agreements) and 81.6 percent within thirty, with a weighted average maturity of four days by June 2026. Even the repo-free 2023 configuration held 53.7 percent within thirty days on the eve of the SVB weekend. Marginal usable stress liquidity $\overline{L}'(S)$ is therefore close to the fund's share of the marginal reserve dollar, and reversing the sign of $\delta_S = \overline{L}' - R'_X$ would require operating commitments to rise by more than two thirds of a dollar per marginal reserve dollar, which no reading of the issuer's disclosed cost base supports. The sign $\delta_S > 0$, which makes both funding flights damaging in Proposition 2, is a measured feature of the reserve technology.

*Three moments of March 2023.* Panel A of Figure 6 shows the three moments that discipline the routing block. First, at the trough on Saturday, March 11, with banking rails closed, USDC traded at 0.874 on Kraken and 0.882 on Binance while USDT/USD held a premium: a 12 to 13 percentage-point discount specific to USDC. The fundamental component was bounded by arithmetic: \$3.3 billion at SVB out of roughly \$40 billion of reserves caps the conditional haircut near 2.5 points even at full resolution probability and a 70 percent uninsured recovery. Through the price block $P = 1 - \sigma - a_B(q_B)$, congestion and route-closure pricing account for at least ten points of the trough. Second, the deposit guarantee announced at 22:15 UTC on Sunday removed the fundamental term while rails stayed closed: the median price moved from 0.971 in the final pre-announcement hour to 0.993 overnight. The residual 0.7-point discount with $\sigma \approx 0$ prices a single object: the closed

Table I: Calibration Inputs from Public Data

| Object | Value | Source |
|---|---|---|
| *Scale* | | |
| USDC in circulation, end-2025 | \$75.27 billion | Filings[a] |
| Circle Reserve Fund net assets, June 2026 | \$61.92 billion (82.3% of circulation) | N-MFP[b] |
| On-chain wrapper assets, October 2025 start | \$13.31 billion | On-chain[c] |
| USDC on the affiliated platform, end-2025 | \$12.50 billion | Filings[a] |
| *Payout and pass-through* | | |
| Fund seven-day gross yield, 2023/24/25/26 means | 4.95 / 5.18 / 4.26 / 3.69% | N-MFP[b] |
| Advertised top-tier USDC rewards, by year | 5.0–5.75 / 5.1–5.2 / 4.1–4.35 / 4.0% | Archive[d] |
| Coinbase agreement costs over reserve income, 2025 | 53.1% | Filings[a] |
| *Reserve ladder* | | |
| Holdings maturing within 7 days, 2025–26 mean | 67.2% (repurchase share 63.0%) | N-MFP[b] |
| Holdings maturing within 30 days, 2025–26 mean | 81.6% (WAM in June 2026: 4 days) | N-MFP[b] |
| *Stress frequency and the fundamental bound* | | |
| Stress windows, 2022–2025 | three, so annual $p \in [0.1, 0.3]$ | On-chain[c] |
| SVB exposure, March 2023 | \$3.3 of ≈\$40 billion reserves; haircut bound $\leq$ 2.5 pp | Filings[a] |
| *March 2023 moments (UTC)* | | |
| Trough with banking rails closed, March 11 | 0.874 (Kraken, 07:15); 0.882 (Binance, 14:00) | Markets[e] |
| USDT/USD, same day | 1.001–1.027 (premium) | Markets[e] |
| Median after the guarantee, rails still closed | 0.993 (final pre-announcement hour: 0.971) | Markets[e] |
| After banks open, March 13 | $\geq$0.995 by 13:50; evening median 0.9992 | Markets[e] |
| Final burns, March 11 versus March 13 | \$1.36 versus \$1.75 billion per day | On-chain[c] |
| Gross burns and net contraction, March 10–17 | \$8.44 and \$6.47 billion | On-chain[c] |
| *Wrapper stress, October 2025* | | |
| Observed stock-loss rate; mark-to-trough | 4.27%; \$1.33 billion (9.96% of start assets) | On-chain[c] |

[a] Circle and Coinbase SEC filings, as documented in Internet Appendix Section IA.IX.
[b] SEC Form N-MFP2/N-MFP3, BlackRock Funds III series S000077205 (the Circle Reserve Fund), monthly, January 2023 to June 2026.
[c] Ethereum ERC-20 event build and stress-window definitions documented in Internet Appendix Section IA.IX; burns are on-chain token retirements and are not classified by economic route.
[d] Wayback Machine captures of platform reward pages; heuristic extraction, region and tier specific.
[e] Kraken public trade tape (minutes rebuilt from trades), Binance Vision, Bitfinex, and Coinbase Exchange public market data.

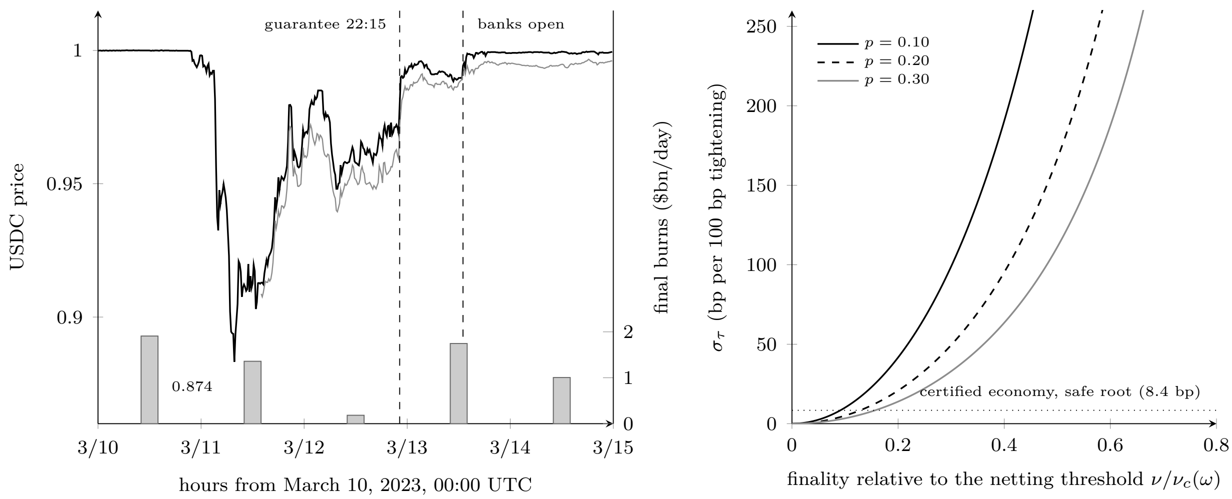


Figure 6: The March 2023 window and the safe-root policy derivative. Panel A plots 15-minute median closes for Kraken USDC/USD (solid) and Binance USDC/USDT (gray) with daily on-chain burns (bars, right axis); dashed lines mark the deposit-guarantee announcement (March 12, 22:15 UTC) and the reopening of U.S. banking rails (March 13, 13:00 UTC). Panel B plots the safe-root derivative $\sigma_\tau$ in the affine benchmark with direct relief $\upsilon$ set to zero, which depends on primitives only through $p$ and $\nu/\nu_c(\omega)$; the dotted line is the certified economy's safe-root value of 8.4 basis points per 100 basis points. Sources: Table I.

redemption route. Third, banks opened at 13:00 UTC on Monday; the price sustained 0.995 within fifty minutes and 0.999 by late afternoon. The decomposition confirms the routing block's premise: when the priority route is rationed, absorption congestion carries stress pricing.

*Quantities and state-contingent finality.* The burn series makes the same point in quantities. Final burns, the on-chain trace of primary-market contraction, were \$1.36 billion on the trough Saturday and \$1.75 billion on reopening Monday: the redemption route settled more volume at a one-point discount than at a twelve-point discount. Quantity, not price, rationed the route while rails were closed, and the platform's weekend suspension of one-for-one conversion is the same fact at the retail boundary. Effective finality collapsed exactly when the outside quote was widest, the state dependence that motivates treating $\nu$ as an architecture variable. Visible order books absorbed about \$1.1 billion the same Saturday, the same order of magnitude as the day's burns, and the week's gross burns of \$8.44 billion against a \$6.47 billion net contraction complete the flow accounting.

*The policy derivative in sufficient-statistic form.* For a design with relative depth $\omega$, write $\nu_c(\omega) := [(1-\Lambda_0)/\{\beta_S p\, d(\omega)\}]^{1/2}$ for its contraction threshold, so the $\nu_c$ of Theorem 2(ii) is

$\nu_c(\frac{1}{2})$. At a stable interior safe root of the affine benchmark with $\upsilon = 0$, which isolates the funding term, (79) reduces to $\sigma_\tau = x^2/[p(1-x^2)]$ with $x := \nu/\nu_c(\omega)$: every primitive enters only through the stress probability and the position relative to the netting threshold. Panel B traces this curve for $p \in \{0.1, 0.2, 0.3\}$. A design a quarter of the way to its threshold converts each 100 basis points of tightening into 22 to 67 basis points of safe-root conditional haircut; halfway, into 111 to 333. The certified economy's safe root, at 8.4 basis points, corresponds to $\nu/\nu_c \approx 0.09$ to 0.16, deep inside the contraction region. The sweep is a local derivative at the safe root, and positive direct relief $\upsilon$ lowers every curve and restores the intended policy sign below the reversal threshold of Theorem 2(ii); the figure quantifies how quickly the funding term alone grows as netting thins.

*Measurement.* Each input the mechanism needs is large at the observed scale: the wrapper layer holds a fifth to a third of circulation, the payout at stake is 370 to 520 basis points, the marginal reserve dollar is overnight liquid, and the one observed stress episode moved the route price and the route quantity in the directions the funding block requires. Sharpening these ranges into point estimates requires the route-level records of Section VI.C: matched tender-to-debit data and two-sided route quotes would deliver $\nu$, $d$, and route tilt, and with them the position of a real architecture in the phase plane of Figure 4. Internet Appendix Section IA.IX documents the construction.

# VIII. Welfare: Capacity, Folds, and Robust Finality

## A. The Real-Resource Ledger

The welfare boundary contains holders, wrapper operators, date-0 outside funding investors, date-1 buyers, authorized redeemers, the issuer, and domestic resolution counterparties with equal weights. Internal payouts, wrapper charges, token-sale expenditures, seller discounts, first-loss allocations, fees, rents, and the incumbent's private franchise value are transfers. Let $\mathcal{U}(W)$ be real service surplus and $\mathcal{C}_H^R(h, W, \sigma, \tau; \kappa_F) := W a_0(h/W, W) + C_F^R(h, \sigma, \tau; \kappa_F)$ actual wrapper inventory and funding cost; it excludes the signed private transfer $\rho_F^A \mathcal{N}$ and, because $C_F^R$ already includes $\zeta_N(\nu)\mathcal{J}$, it includes the facility's collateral, uncovered-exposure, and operating resources. If direct-access eligibility requires real investment $\Gamma_A(\varrho_O)$, the expected access cost not already in direct-redemption processing is $\mathcal{C}_A^R := \Gamma_A(\varrho_O) + p\mathcal{N}\varrho_O \int_0^{[\sigma-\varphi_D]_+} c \, \mathrm{d}F_O(c)$, with the access fee a transfer and exercise cost real; the baseline sets $\mathcal{C}_A^R = 0$. The withdrawal cutoff implies the real timing loss $\mathcal{M}(D; W) := W t_H(D)$, because flexible holders exit only to avoid dilution. Original-holder redemption, market absorption, authorized processing, and emergency finance have real costs $C_D(\widehat{Z}; W)$, $A_B(q_B)$, $C_R(Q)$, and $F_I(q)$; $F_I$ excludes reserve destruction and resolution resources counted through $\Theta\pi(q)$; date-0 redemption and par secondary sale of backing are internal transfers; $C_X(S)$ is any real resource use behind the operating load $R_X(S)$. In resolution, decompose the missing value per junior token as $\alpha = \chi_A + \chi_X + \chi_T$, where $\chi_A$ is destroyed value or resolution cost, $\chi_X$ leaves the boundary, and $\chi_T$ is internal, and write $\chi_J := \chi_A + \chi_X \in [0, \alpha]$; global surplus sets $\chi_X = 0$, not mechanically $\chi_J = \alpha$. If $D_R$ is operational disruption per resolution, social value at risk is $\Theta = D_R + \chi_J N_J$, and the date-0 real cost net of service surplus is

$$\begin{aligned} \mathcal{K} = & -\mathcal{U}(W) + \mathcal{C}_H^R(h, W, \sigma, \tau; \kappa_F) + \mathcal{C}_A^R + g_I(\ell) \\ & + p\{W\phi(x/W) + \mathcal{M}(D; W) + C_D(\widehat{Z}; W) \\ & \qquad + A_B(x - Q) + C_R(Q) + C_X(S) + F_I(q) + \Theta\pi(q)\}. \end{aligned} \tag{83}$$

Adding seller discounts, late-holder dilution, or $\sigma N_J$ would double count private incidence as a second resource loss.

## B. The Architecture Loss and the Recovery Shadow

The settlement-design block is a constrained projection of (83), not a second criterion. Holding nonarchitecture terms fixed, let $\mathcal{L}_B(\sigma, j, \nu)$ denote the remaining real resolution loss at issuer-facing tender $j$, twice continuously differentiable with no direct depth argument and $\mathcal{L}_{B,\sigma} \geq 0$. Because $\sigma = \alpha\pi$, the ledger gives the exact embedding $\mathcal{L}_B(\sigma, j, \nu) = (p/\alpha)\Theta_B(j, \nu)\sigma$ with $\Theta_B(j, \nu) = D_R + \chi_J N_J^B(j, \nu)$. In the transparent stock-retirement projection, stress settlements other than the architecture debit are held fixed, so $N_J^B(j, \nu) = \overline{N} - \nu j$, $\mathcal{L}_{B,j} = -p\chi_J \nu\pi$, and $\mathcal{L}_{B,\sigma} = p\Theta_B/\alpha$: capacity that raises issuer-facing tender can raise resolution probability while retiring claims exposed to loss in resolution. The constant-value-at-risk benchmark sets $\chi_J = 0$ (or holds $\Theta_B = \bar{\Theta}$), giving

$$\mathcal{L}_B(\sigma, j, \nu) = \Upsilon\sigma, \qquad \Upsilon := \frac{p\bar{\Theta}}{\alpha} > 0. \tag{84}$$

On a stated smooth stable selection, a planner choosing committed instruments $z$ minimizes $\mathcal{K}(\sigma, z) + \Gamma(z)$ subject to $\sigma = \mathcal{T}(\sigma, z)$; with the recovery shadow

$$\mu_R := \frac{\mathcal{K}_\sigma}{1 - \mathcal{T}_\sigma}, \qquad \frac{\mathrm{d}\sigma}{\mathrm{d}z_j} = \frac{\mathcal{T}_{z_j}}{1 - \mathcal{T}_\sigma}, \qquad \frac{\mathrm{d}\mathcal{K}}{\mathrm{d}z_j} = \mathcal{K}_{z_j} + \mu_R \mathcal{T}_{z_j}, \tag{85}$$

an interior optimum satisfies $\mathcal{K}_{z_j} + \mu_R \mathcal{T}_{z_j} + \Gamma_{z_j} = 0$.

## C. The Capacity Externality and Its Fold Amplification

The wrapper sector purchases proportional route depth $d > 0$ at price $q_d$ from competitive providers with real cost $\Gamma_F(d;\xi)$, $\Gamma_F(0;\xi) = 0$, $\Gamma_{F,d} > 0$, $\Gamma_{F,dd} \geq 0$, who supply until $q_d = \Gamma_{F,d}$; the regulator adds $\Gamma_F$ to the ledger and internalizes recovery after this conditional market clears.

Providers of depth sell cheaper reallocation between the two funding routes; buyers internalize the cost saving but not the fact that cheaper reallocation amplifies whichever route tilt prevails and moves system recovery. When the tilt runs toward the outside route ($r_N < 0$, low finality), extra depth pulls funding *away* from redemption, improves recovery, and the market underprovides depth; when the tilt runs toward the issuer ($r_N > 0$, high finality), the externality reverses and the market overprovides. The correcting charge is proportional to the same net-curvature margin $\mathcal{M}_F$ that measures distance to a fold, so its magnitude grows as the branch approaches the fold.

Theorem 3 (Capacity distortion near a recovery fold): *Retain an interior allocation, $\rho_F \geq 0$, the architecture loss $\mathcal{L}_B$, and the objects of Theorem 1. The robust statements hold on a reached domain with $\mathcal{M}_F > 0$; the singular limits approach a generic fold satisfying* (64).

(i) The exact capacity wedge and its implementing charge. *On any robust domain with $\mathcal{M}_F > 0$ throughout, recovery and the funding split are unique, and*

$$\operatorname{sgn} r_N(\nu, d) = \operatorname{sgn} \Delta_N^G(\nu) \quad \textit{for every admissible } d: \tag{86}$$

*depth changes the magnitude of the funding tilt but not its direction. Its equilibrium responses are*

$$\mathcal{J}_d = \frac{r_N}{d\,\mathcal{M}_F}, \qquad \sigma_d^B = \frac{\beta_S \nu}{\mathcal{R}_{B,\sigma}} \frac{r_N}{d\,\mathcal{M}_F}. \tag{87}$$

*The conditional capacity market clears at the zero of*

$$\mathcal{F}_{P,d} = \Gamma_{F,d}(d;\xi) - \frac{\Phi_F(\mathcal{J};h,\omega)}{d^2}. \tag{88}$$

*Define the marginal social exposure loading* $\Omega_B := \nu\rho_F + \mathcal{L}_{B,j} + (\beta_S\nu/\mathcal{R}_{B,\sigma})\,\mathcal{L}_{B,\sigma}$*. The social residual, including the real funding ledger and the loss inherited from* (83)*, obeys the exact general wedge, and a differentiable capacity charge* $\mathcal{P}(d)$ *added to the provider's objective implements the planner's residual pointwise on a selected smooth stable branch, when*

$$\mathcal{P}'(d) = \mathcal{F}_{S,d} - \mathcal{F}_{P,d} = \Omega_B \frac{r_N}{d\,\mathcal{M}_F}, \tag{89}$$

*unique up to an additive constant if residual equality is required throughout the branch; a positive derivative is a marginal charge, a negative one a subsidy. Internet Appendix Section IA.VII records equivalent forms of this wedge: the constant-value-at-risk reduction, a stability-margin form proportional to* $1/(1-\Lambda_F^G)$*, and the stock-retirement loading* (IA.207)*, whose positivity condition* (IA.208) *makes the capacity distortion carry the sign of* $\Delta_N^G$*: automatic under constant value at risk for* $\nu > 0$*, but possibly reversed when retiring junior claims saves substantial real resolution loss. If* $\Omega_B > 0$ *and both residuals are strictly increasing with single crossings, the private and social depths satisfy*

$$d_S = d_P \text{ if } \nu \in \{0, \nu_A\}, \qquad d_S > d_P \text{ if } 0 < \nu < \nu_A, \qquad d_S < d_P \text{ if } \nu > \nu_A. \tag{90}$$

*Approaching a fold of Theorem 1 along its stable branch at fixed depth, the net curvature vanishes at the square-root rate* $\mathcal{M}_F \sim \sqrt{2D_f|C_f|\,|\nu-\nu_f|}$*, so if* $r_N^f \neq 0$*,*

$$|\mathcal{J}_d|\,,\ \left|\sigma_d^B\right| = \Theta(|\nu-\nu_f|^{-1/2}), \tag{91}$$

*and if in addition* $\Omega_B^f \neq 0$*,*

$$|\mathcal{F}_{S,d} - \mathcal{F}_{P,d}| = \Theta(|\nu-\nu_f|^{-1/2}). \tag{92}$$

*The recovery fold is thus also a capacity-shadow singularity: unless the route tilt or welfare exposure is tangent* ($r_N^f = 0$ *or* $\Omega_B^f = 0$)*,* $\mathcal{P}'$ *diverges, no bounded marginal capacity schedule implements the planner uniformly near the fold, and a schedule contingent only on* $(\nu, d)$ *cannot implement distinct branch wedges at one design point; the recovery-selection institution* *must be specified.*

(ii) The route-reversal threshold. *If $\Delta_N^G(0) < 0 < \Delta_N^G(1)$ and $(\Delta_N^G)'(\nu) > 0$, the tilt has a unique zero $\nu_A \in (0,1)$: there additional depth changes neither tender nor recovery, and private and social depth coincide. When $\Omega_B > 0$ it is also the unique capacity-distortion threshold; if $\Omega_B$ changes sign, the welfare boundary creates an additional reversal even though the funding route does not change direction. For the affine benchmark these single-crossing conditions follow from $\zeta_N(0) > 0$, $\Delta_N^A(1) > 0$, and the positive-drift condition (IA.221) in the Internet Appendix. Writing $\mathcal{R}_{B,x}^A := \mathcal{R}_{B,x}(\sigma_0(\nu_A); \tau)$, settlement costs raise $\nu_A$, issuer-facing baseline share lowers it, and when $\Omega_B > 0$ throughout the branch a tighter yield restriction makes the capacity market switch from under- to overprovision at lower finality whenever $\mathcal{R}_{B,\sigma}^A > p\mathcal{R}_{B,\tau}^A$; the affine benchmark reduces this to $1 - \Lambda_0 > p\upsilon$. If the selected stable branch reaches the fold with positive net curvature between its allocation and $\mathcal{J}_0$ before the limit,*

$$\operatorname{sgn}(\nu_f - \nu_A) = \operatorname{sgn} \Delta_N^G(\nu_f) = \operatorname{sgn} r_N^f : \tag{93}$$

*a positive sign means competitive capacity becomes excessive before the stable branch reaches the fold and the positive distortion then diverges; a negative sign means the fold arrives while capacity remains underprovided.*

When the route-tilt and welfare-exposure numerators remain nonzero, the implementing marginal capacity charge diverges as a stable branch approaches a generic fold. On a robust branch the Pigouvian schedule (89) is exact; near a fold, where branches coexist, no fee contingent on observable design variables prices both at once, and the design problem passes to architectural instruments, including finality caps and netting.

## D. Robust Finality

How much finality should the access contract itself permit? The trade-off is real resources against recovery pressure: netting (lower $\nu$) consumes settlement credit through collateral carry and uncovered-exposure cost, summarized by $\zeta_N$, while finality (higher $\nu$) completes more pre-stress redemption and raises the haircut. Working on a fixed-depth robust domain that caps the loop gain away from one, the trade-off has a clean interior solution.

Fix $h > 0$, $\omega \in (0,1)$, $\tau$, and write $d := d(\omega) > 0$. Admit only architectures with loop gain at most $\overline{\Lambda} \in (\Lambda_0, 1)$, that is, $\nu \in [0, \overline{\nu}(\omega)]$ with $\overline{\nu}(\omega) := \min\left\{1, \sqrt{(\overline{\Lambda} - \Lambda_0)/[\beta_S p\, d(\omega)]}\right\}$, with the projected fixed point and funding split interior throughout. Then $\mathcal{D}_B(\nu) :=$

$1-\Lambda_0-\beta_S p\nu^2 d > 0$, the unique globally stable recovery and quote are $\sigma^B(\nu) = A_{\omega,\nu}(\tau)/\mathcal{D}_B(\nu)$ and $\rho_F(\nu) := p\sigma^B(\nu) - i(\tau)$, and with $Y(\nu) := \nu\rho_F(\nu)$ and $B_N(\nu) := \omega h - d\,\zeta_N(\nu)$, the quadratic allocation and settlement tariff give integrated real funding and settlement cost $C_F^{R,B}(\nu) = \kappa_0 h^2/2 + \omega h\,\zeta_N(\nu) + (d/2)\{Y(\nu)^2 - \zeta_N(\nu)^2\}$, and, under constant value at risk, the regulator minimizes $\mathcal{K}_N^B(\nu) := C_F^{R,B}(\nu) + \Upsilon\sigma^B(\nu)$ over $[0, \overline{\nu}(\omega)]$.

Proposition 3 (Robust interior finality): *Suppose $\rho_F(\nu) \geq 0$ and $q_A f_A \geq \rho_F(\nu) + m_N(\nu)$ throughout the robust domain, and the curvature conditions $g_N \leq 2c_N$, $B_N(\nu) > 0$, and $g_N B_N(\nu) \geq d\,m_N(\nu)^2$ hold on $[0, \overline{\nu}(\omega)]$. Then $\sigma^B$ is strictly convex and increasing in $\nu$ with $\sigma_\nu^B = \beta_S[\mathcal{J}(\nu) + \nu d\{\rho_F(\nu) + m_N(\nu)\}]/\mathcal{D}_B(\nu) > 0$, the objective $\mathcal{K}_N^B$ is strictly convex, and the robust finality choice $\nu^R$ is unique. It is interior if and only if*

$$\Upsilon\frac{\beta_S}{1-\Lambda_0} < c_N + g_N \quad \textit{and} \quad dY(\overline{\nu})Y_\nu(\overline{\nu}) + \Upsilon\sigma_\nu^B(\overline{\nu}) > m_N(\overline{\nu})B_N(\overline{\nu}), \tag{94}$$

*in which case it uniquely solves*

$$m_N(\nu^R)B_N(\nu^R) = dY(\nu^R)Y_\nu(\nu^R) + \Upsilon\sigma_\nu^B(\nu^R), \qquad Y_\nu = \rho_F + \nu p\sigma_\nu^B. \tag{95}$$

*Full netting is optimal, $\nu^R = 0$, if and only if $\Upsilon\beta_S/(1-\Lambda_0) \geq c_N + g_N$; the cap binds, $\nu^R = \overline{\nu}(\omega)$, if and only if the second inequality in (94) fails. At an interior solution, $\mathrm{d}\nu^R/\mathrm{d}\Upsilon < 0$, and the sign of $\mathrm{d}\nu^R/\mathrm{d}x$ for settlement costs $x \in \{c_N, g_N\}$ compares the direct saving from finality against the endogenous route-substitution and recovery response (exact threshold in the Internet Appendix). If the cap binds with $\overline{\nu}(\omega) < 1$, the loop-gain constraint mandates at least $1 - \overline{\nu}(\omega) > 0$ netting.*

The left side of the first-order condition is the marginal settlement-credit saving from one more unit of finality, weighted by the netting-exposed base; the right side is the marginal funding-cost and social-recovery pressure that finality creates. Social value at risk therefore maps monotonically into netting: systems whose resolution is more disruptive should complete less pre-stress redemption, with full netting exactly when $\Upsilon\beta_S/(1-\Lambda_0) \geq c_N + g_N$. The proposition is stated on a domain that caps the loop gain, which is what "robust" means here: it is the design problem a regulator faces *after* imposing the stability margin that Theorem 3 shows price instruments cannot deliver near a fold. The settlement-credit microfoundation and proof are Internet Appendix Proposition IA.1.

# IX. Conclusion

A yield restriction on stablecoin issuers does not extinguish the demand for remunerated money; it relocates the promise of liquidity to wrappers whose backing remains a claim on the restricted issuer. The paper's contribution is to model the connection that the Regulation Q episode lacked: the route by which the wrapper funds its liquidity, final redemption or outside placement, is priced by investors who receive the restricted payout and bear junior losses, so both beliefs and the restriction itself move the funding of settlement. Finality's asymmetry, linear in exposure and quadratic in feedback, organizes the funding results. In embedded fixed-cash, fixed-claim-scale benchmarks, it yields the quartic-congestion cusp in feedback leverage and route tilt and the affine–quadratic netting threshold $\nu_c$. In the full contract, finality enters the four-force decomposition of recovery amplification and yield policy.

The certified full-contract economy exhibits a clean division of labor. Routing and claims create the unstable middle crossing and hence multiplicity, while policy-induced funding flight alone flips the yield-policy sign at the stable safe root. A model of the claim without its funding chain can therefore count equilibria correctly and still misjudge the restriction. When the relevant route-tilt and welfare-exposure numerators remain nonzero, the price-based capacity correction diverges near a generic fold; netting, finality caps, and the selection institution remain available as architectural instruments.

With exogenous recovery-news variation, matched tender–debit records, and route quotes, feedback leverage and route tilt would be estimable under the benchmark specification. Testing hysteresis would additionally require repeated directional changes in the same architecture. Separately, the full-contract decomposition predicts policy effects even at zero claim migration. The calibrated illustration pins down levels, the reserve ladder, and the March 2023 moments; route-level records would complete the measurement. If stablecoin regulation repeats the Regulation Q experiment, the restriction's effect on recovery will be decided in the funding market behind the wrapper, which a claim-only model does not observe.

# References

Ahmed, Rashad, Iñaki Aldasoro, and Chanelle Duley, 2024, Public information and stablecoin runs, BIS Working Paper 1164, Bank for International Settlements, Revised January 2025.

Anadu, Kenechukwu E., Pablo D. Azar, Marco Cipriani, Thomas M. Eisenbach, Catherine Huang, Mattia Landoni, Gabriele La Spada, Marco Macchiavelli, Antoine Malfroy-Camine, and J. Christina Wang, 2023, Runs and flights to safety: Are stablecoins the new money market funds?, Research Department Working Paper 23-11, Federal Reserve Bank of Boston, Revised March 2024.

Balasubramaniam, Swaminathan, 2026, Scale, yield, and run risk in yield-bearing stablecoins, Working paper, NEOMA Business School, SSRN 6383598.

Bank of England, 2026, Sterling-denominated systemic stablecoins, Policy statement and consultation on draft Code of Practice, Bank of England, Published June 22, 2026; accessed July 18, 2026.

Barrios, John Manuel, Christoph Bertsch, and Linda Schilling, 2026, One asset, two financial systems: Stablecoins and the transmission of runs between decentralized and traditional finance, Working Paper 6963621, SSRN.

Brunnermeier, Markus K., and Yuliy Sannikov, 2014, A macroeconomic model with a financial sector, *American Economic Review* 104, 379–421.

Capponi, Agostino, and Jin-Wook Chang, 2025, Settlement speed and financial stability, Finance and Economics Discussion Series 2025-101, Board of Governors of the Federal Reserve System, Revised June 5, 2026.

Carapella, Francesca, Arazi Lubis, and Alexandros Vardoulakis, 2026, Stablecoins in 2025: Developments and financial stability implications, Feds notes, Board of Governors of the Federal Reserve System.

Circle, 2023, $3.3 billion of USDC reserve risk removed, dollar de-peg closes, Press release, Published March 13, 2023; press release dated March 12, 2023.

Circle, 2025, USDC terms, Issuer terms and conditions, Last updated December 12, 2025; accessed July 18, 2026.

Circle, n.d., How minting and redemption works, Circle Mint developer documentation, Accessed July 18, 2026.

Circle Internet Group, Inc., 2026, Annual report for the fiscal year ended december 31, 2025, Form 10-K filed with the U.S. Securities and Exchange Commission, Filed March 9, 2026; SEC accession 0001876042-26-000062; accessed July 18, 2026.

Committee on Payments and Market Infrastructures, and Board of the International Organization of Securities Commissions, 2022, Application of the principles for financial market infrastructures to stablecoin arrangements, CPMI Papers, Bank for International Settlements and International Organization of Securities Commissions, Published July 13, 2022.

Council of Economic Advisers, 2026, Effects of stablecoin yield prohibition on bank lending, Technical report, Executive Office of the President of the United States.

d'Avernas, Adrien, Vincent Maurin, and Quentin Vandeweyer, 2026, Can stablecoins be stable?, *Management Science* Articles in Advance, published online May 26, 2026.

Diamond, Douglas W., and Philip H. Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91, 401–419.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2017, The deposits channel of monetary policy, *Quarterly Journal of Economics* 132, 1819–1876.

European Systemic Risk Board, 2023, Crypto-assets and decentralised finance: Systemic implications and policy options, Report, European Systemic Risk Board.

Galbiati, Marco, and Kimmo Soramäki, 2010, Liquidity-saving mechanisms and bank behaviour, Working Paper 400, Bank of England.

Garcia Ocampo, Denise, 2025, Stablecoin-related yields: Some regulatory approaches, FSI Briefs 27, Bank for International Settlements.

Gilbert, R. Alton, 1986, Requiem for regulation Q: What it did and why it passed away, *Federal Reserve Bank of St. Louis Review* 68, 22–37.

Goel, Tirupam, Ulf Lewrick, and Isha Agarwal, 2026, Making stablecoins stable(r): Can regulation help?, BIS Working Paper 1355, Bank for International Settlements.

Goldstein, Itay, and Ady Pauzner, 2005, Demand–deposit contracts and the probability of bank runs, *Journal of Finance* 60, 1293–1327.

Gorton, Gary, and Andrew Metrick, 2012, Securitized banking and the run on repo, *Journal of Financial Economics* 104, 425–451.

Gorton, Gary B., Elizabeth C. Klee, Chase P. Ross, Sharon Y. Ross, and Alexandros P. Vardoulakis, 2026, Leverage and stablecoin pegs, *Journal of Financial and Quantitative Analysis* 61, 99–136.

Gross, Marco, and Richard Senner, 2026, From par to pressure: Liquidity, redemptions, and fire sales with a systemic stablecoin, IMF Working Paper 2026/005, International Monetary Fund.

Huang, Wenqian, Nikola Tarashev, and Xinyi Wang, 2026, Stablecoin remuneration on centralised exchanges, BIS Bulletin 125, Bank for International Settlements.

Kacperczyk, Marcin, and Philipp Schnabl, 2013, How safe are money market funds?, *Quarterly Journal of Economics* 128, 1073–1122.

Kahn, Charles M., and William Roberds, 1998, Payment system settlement and bank incentives, *Review of Financial Studies* 11, 845–870.

Klee, Elizabeth C., Arazi Lubis, Chase P. Ross, Sharon Y. Ross, and Alexandros P. Vardoulakis, 2026, The fragility of perfectly safe digital money, Finance and Economics Discussion Series 2026-037, Board of Governors of the Federal Reserve System.

Koh, Pairie, and Kuntara Pukthuanthong, 2026, Wrapped stablecoins and the DeFi systemic risk, Working paper, February 12, 2026.

Lee, Michael Junho, 2026, Composable finance, Staff Report 1177, Federal Reserve Bank of New York.

Lee, Michael Junho, and Donny Tou, 2026, Stablecoin disintermediation, Staff Reports 1185, Federal Reserve Bank of New York.

Li, Bo, Tommaso Mancini-Griffoli, Marcello Miccoli, Brandon Joel Tan, and Longmei Zhang, 2026, Making stablecoins stable, IMF Working Paper 2026/074, International Monetary Fund.

Li, Ye, and Simon Mayer, 2025, Financial intermediation with fragile seigniorage: A dynamic model of stablecoin issuers, Fisher College of Business Working Paper 2020-03-030 and Charles A. Dice Center Working Paper 2020-30; revised February 5, 2026.

Lyons, Richard K., and Ganesh Viswanath-Natraj, 2023, What keeps stablecoins stable?, *Journal of International Money and Finance* 131, 102777.

Ma, Yiming, Yao Zeng, and Anthony Lee Zhang, 2025, Stablecoin runs and the centralization of arbitrage, NBER Working Paper 33882, National Bureau of Economic Research, Forthcoming, Review of Financial Studies.

Matta, Rafael, and Enrico Perotti, 2024, Pay, stay, or delay? How to settle a run, *Review of Financial Studies* 37, 1368–1407.

Office of the Comptroller of the Currency, U.S. Department of the Treasury, 2026, Implementing the guiding and establishing national innovation for U.S. stablecoins act for the issuance of stablecoins by entities subject to the jurisdiction of the office of the comptroller of the currency, *Federal Register* 91, 10202–10303, Notice of proposed rulemaking; Docket ID OCC-2025-0372; RIN 1557-AF41; FR Doc. 2026-04089.

Pagès, Henri, and David B. Humphrey, 2005, Settlement finality as a public good in large-value payment systems, Working Paper Series 506, European Central Bank.

Pan, Kevin, and Yao Zeng, 2017, ETF arbitrage under liquidity mismatch, Working Paper Series 59, European Systemic Risk Board.

Parlour, Christine A., Uday Rajan, and Johan Walden, 2022, Payment system externalities, *Journal of Finance* 77, 1019–1053.

Pennacchi, George, 2012, Narrow banking, *Annual Review of Financial Economics* 4, 141–159.

Plantin, Guillaume, 2015, Shadow banking and bank capital regulation, *Review of Financial Studies* 28, 146–175.

Schmidt, Lawrence, Allan Timmermann, and Russ Wermers, 2016, Runs on money market mutual funds, *American Economic Review* 106, 2625–2657.

Team Circle, 2023, March 15, 2023 | update on USDC operations, Company update.

Tierno, Paul, and Marc Labonte, 2026, The stablecoin yield debate, In Focus IF13174, Congressional Research Service, Released March 6, 2026.

United States Congress, 2025, Guiding and Establishing National Innovation for U.S. Stablecoins Act, Public Law 119–27, Approved July 18, 2025; 139 Stat. 419.

U.S. Department of the Treasury, Board of Governors of the Federal Reserve System, and Federal Deposit Insurance Corporation, 2023, Joint statement by the department of the treasury, federal reserve, and FDIC, Joint release dated March 12, 2023.

Watsky, Cy, Jeffrey Allen, Hamzah Daud, Jochen Demuth, Daniel Little, Megan Rodden, and Amber Seira, 2024, Primary and secondary markets for stablecoins, FEDS Notes, Board of Governors of the Federal Reserve System.

Yuan, Kathy Zhichao, 2026, Market-based liquidity transformation, Working paper, May 20, 2026, SSRN 6806178.