PUBLIC SAMPLE — REDACTED. This is a genuine Vari review with the client's identity and identifying scope details removed (shown as [redacted]). Methodology, findings, severities, and recommendations are unaltered — shared as a representative example of Vari's full-stack review work.

VARI

ECONOMIC MODELING ADDENDUM · REV 2

ECONOMIC & SOLVENCY MODELING

[redacted] — [redacted] & [redacted] + Omnichain Stablecoin

CLIENT [redacted]	REPORT DATE 2026-08-11
COMPANION TO Full-Stack Security Review	METHOD Quantitative red-team
LIVE SUPPLY MODELED S ≈ $[redacted]M	PARAM SET 6.5d vest · 14d cooldown · fees on

00 PURPOSE & METHOD

The main review establishes that [redacted] has **no critical code vulnerability and no permissionless path to mint unbacked tokens, drain collateral, or double-spend across chains**. This addendum goes past the module boundary that a contract audit stops at, and asks a different question: **under what market and behavioral conditions does a fully-honest, un-exploited [redacted] still break its peg or fail to honor redemptions?** These are economic and mechanism-design failure modes — bank runs, reflexive depegs, collateral-quality drift, tranche-run dynamics, cross-chain peg fragmentation — invisible to line-by-line review because the code behaves exactly as written; it is the *incentive geometry around the code* that fails.

Every figure below is generated from a parameterized model of the live system at the **current configured parameters**: single hub supply S ≈ $[redacted]M, PayFi book duration ~45 days, **6.5-day reward vesting**, **14-day unstake cooldown**, **mint/redeem fees enabled**, no oracle in the mint/redeem path, and a per-user redemption queue whose ordering is **FIFO as replayed by the protocol indexer**. The models are deliberately **structural, not predictive**: they identify the *shape* of the risk surface and the handful of governance parameters that move an operator from the safe region to the unsafe one. Off-chain quantities we cannot observe (actual liquid buffer, PayFi maturity ladder, insurance-fund size) are swept across ranges, with the region a prudent operator should stay inside marked.

> **What this revision reflects.** Three parameter changes and one architectural clarification since Rev 1 have **closed or bounded three of the four original existential vectors**: the **14-day cooldown** doubles the loss-recognition window a run must beat (junior-tranche run); **live mint/redeem fees** neutralize the small-depeg collateral-drain arbitrage and act as a run brake (Gresham drift); and the **indexer-enforced FIFO queue** removes the informational asymmetry that drove the reflexive-depeg spiral. What none of them touches is the **liquid buffer ratio r** — so this revision sharpens rather than softens the headline: **r is now the lone remaining existential variable, and enforcing a floor on it is the single highest-leverage economic hardening left.**

1	8	6	7
EXISTENTIAL (S-TIER)	SEVERE (A-TIER)	BOUNDED (B-TIER)	QUANT MODELS

01 THE TWO MASTER VARIABLES

[redacted] is a maturity-transformation machine: it issues an instantly-redeemable dollar ([redacted]) and a savings share ([redacted]+) against a book of **short-but-not-instant** PayFi credit that pays >10% precisely *because* it is illiquid. That is the same structure as a bank, and it inherits the same failure mode — a redemption demand that exceeds the liquid buffer before the credit book matures. Two numbers decide whether that ever happens.

R — THE LIQUID BUFFER RATIO

Let `r` be the fraction of supply `S` held as instantly-redeemable on-chain reserve (USDC/USDT in the silos), with the remaining `(1−r)` deployed into PayFi credit of average maturity `D ≈ 45 days`. Yield and safety pull `r` in opposite directions: hitting a >10% APR target on a book that yields, say, ~13% gross requires deploying the large majority of assets, which forces `r` low — plausibly into the **5–20%** band. The instant-redemption promise is only as deep as `r·S`, and — critically — **none of the Rev-2 parameter changes add a single dollar to it.** Fees, a longer cooldown, and a fair queue make a run slower, costlier, and more orderly; they do not create liquidity that has not matured.

```
instant redemption capacity = r · S   |   queue-bound (illiquid) = (1 − r) · S   |
   daily maturity inflow ≈ (1 − r)·S / D
```

Δ — THE LOSS-RECOGNITION LAG

When a PayFi obligor defaults or a receivable is impaired, the loss is real off-chain immediately, but it only reaches the on-chain share price when an operator calls `burnAssets` (which also auto-pauses the vault). The interval Δ between those two events is a window in which **the share price is knowingly stale** — [redacted]+ trades and redeems at a price that overstates backing. The **14-day cooldown widens the Δ the system can tolerate to 14 days**: an informed staker can no longer exit inside the lag unless the lag exceeds the cooldown (§06). Δ remains the variable that converts an ordinary credit loss into a run, but it is now defused by policy rather than being a knife-edge at one week.

> **Neither r nor Δ is bounded on-chain.** There is no invariant forcing a minimum buffer, and no mechanism capping how long a known loss can sit un-booked. The 14-day cooldown makes the Δ requirement forgiving, but `r` is still entirely operator discipline. The single highest-leverage economic hardening [redacted] can do is to **publish and enforce a minimum-r policy, and disclose a maximum-Δ policy comfortably inside 14 days**.

02 LIQUIDITY COVERAGE & DAYS-OF-COVER

The first-order question for any redeemable-on-demand instrument is: **how much can be redeemed instantly, and for how many days can the buffer absorb a sustained outflow before the queue takes over?** The left panel splits supply into the instantly-redeemable buffer ($r \cdot S$) and the queue-bound remainder. The right panel plots days-of-cover — how long the buffer lasts once PayFi maturities (inflow) are outpaced by redemptions (outflow) — as a function of r, for three outflow regimes.

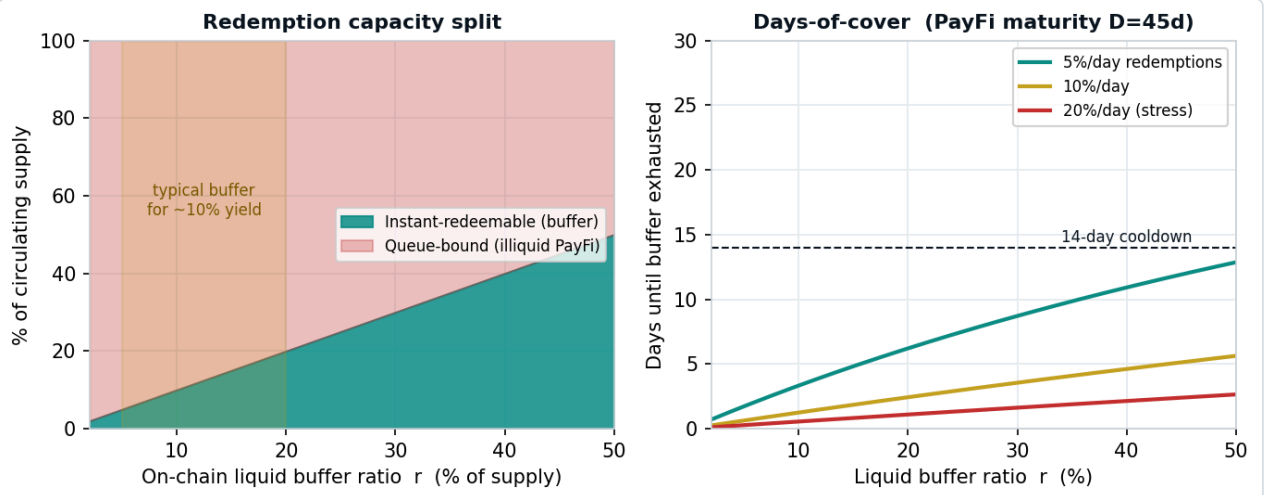

Fig. 1 — Redemption capacity and days-of-cover. Left: at a typical yield-seeking buffer ($r \approx 5$–20%, amber band), only ~5–20% of supply is instantly redeemable; the rest is queue-bound behind illiquid PayFi credit. Right: under a 20%/day stress outflow, even a 20% buffer is exhausted in under three days — far inside the 14-day cooldown that [redacted]+ stakers cannot exit within. Days-of-cover collapses non-linearly as r falls.

The operational reading: the 14-day unstake cooldown on [redacted]+ is a **defense for the protocol, not the staker** — it holds savings-share (junior) assets in place for two weeks while a run plays out, which is longer than the buffer survives any real stress and therefore a genuinely useful circuit-breaker. But [redacted] itself has **no cooldown**; its holders are first in the queue and their instant-redemption depth is exactly $r \cdot S$. If r is run at the low end to maximize yield, the buffer funds only a small fraction of supply, and any coordinated redemption larger than the buffer immediately routes into the queue — modeled in §04.

03 BANK-RUN PROBABILITY (MONTE CARLO)

Days-of-cover is deterministic; real redemption demand is stochastic and fat-tailed. We simulate 20,000 seven-day paths of aggregate redemption demand under three behavioral regimes — calm, nervous, and crisis — each a base withdrawal rate plus a Poisson-timed lognormal "spike" process that models coordinated exits triggered by news, a competitor depeg, or a social-media cascade. For each buffer ratio r we compute the probability that cumulative 7-day demand exceeds the redeemable capacity (buffer plus one week of maturities). This models the **senior [redacted] dollar**, which has no cooldown; the 14-day junior cooldown does not enter here.

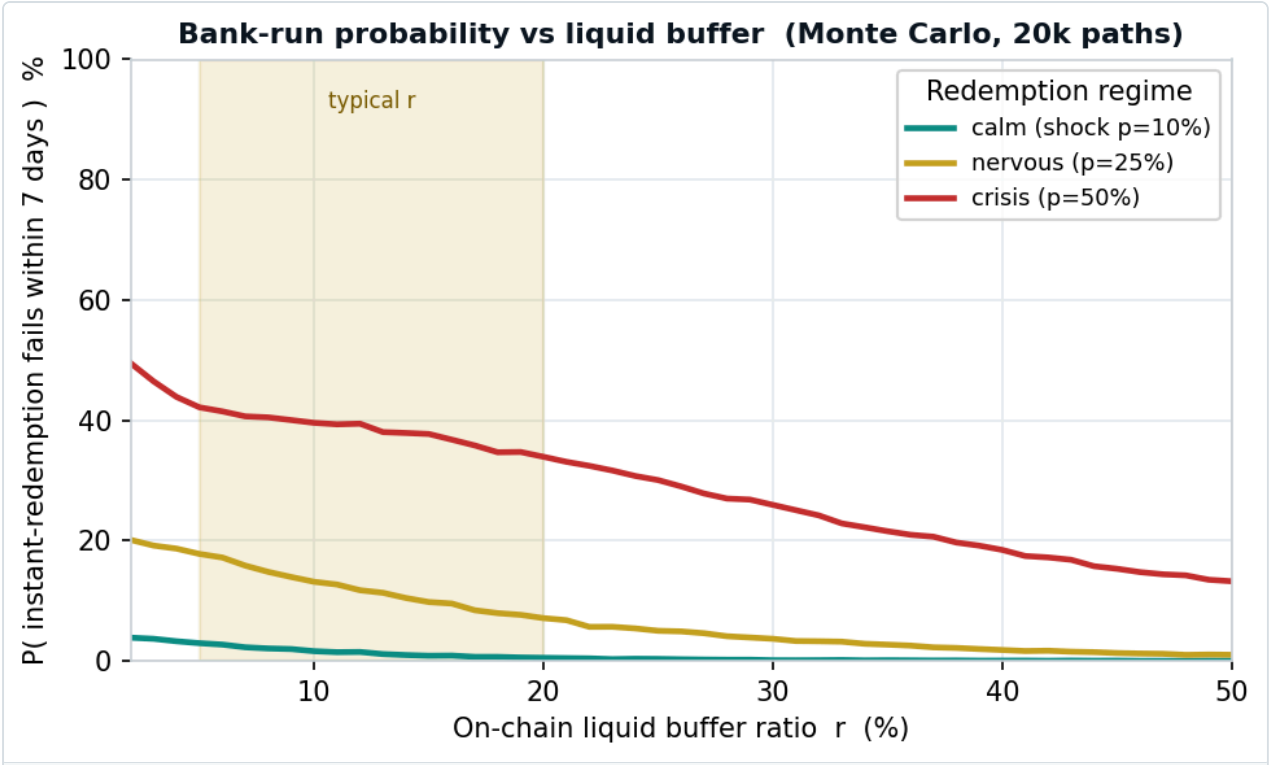

Fig. 2 — Probability instant redemption fails within 7 days. 20,000 Monte-Carlo paths per point. In the calm regime the buffer is comfortable across the whole plausible-r band. In the crisis regime (coordinated exit probability 50%), a protocol running r ≈ 10% faces a **double-digit-to-majority** probability of exhausting instant liquidity inside a week — despite being fully solvent on a mark-to-maturity basis. Insolvency is not required for a run; illiquidity is sufficient.

Diamond–Dybvig applies, now with a partial brake. Redemption honored at par gives a rational holder who assigns any probability to others running a first-mover incentive to run first — the classic multiple-equilibrium structure. The buffer ratio `r` sets how large a shock is needed to tip from the "everyone stays" to the "everyone runs" equilibrium. A **redeem fee (§05) raises that tipping threshold** by taxing the exit, and the **indexer FIFO queue (§04) removes the panic-amplifier**, but neither changes the underlying solvency-vs-liquidity gap — a higher `r` is still the only thing that raises the shock the buffer can actually absorb.

04 THE REDEMPTION QUEUE — INDEXER-ENFORCED FIFO

In Rev 1 this was the single most severe [redacted]-specific vector, because the queue combined three adverse properties: per-user (no global order), discretionary completion by the solver, and opaque depth. **The FIFO clarification changes the picture materially.** While ordering is not enforced by the token contract, the protocol indexer provides an **accurate, replayable FIFO reconstruction of the queue** from on-chain request events — every holder can see the line and their position in it. That removes the property that made the vector existential: the **informational asymmetry** that let a panicking holder assume they were behind everyone else and therefore rush to redeem first.

The reflexive-depeg spiral in Rev 1 ran: opaque queue → holders price redemption risk on rumor → run to be safe → longer queue → wider secondary discount → more running. A **legible, ordered queue breaks the loop at the first link** — a holder who can see they are near the front has no reason to panic-sell into a discount, and one who is far back learns that from the queue rather than from a collapsing price. The phase plane below is therefore now best read as a **liquidity boundary, not a panic boundary**: crossing it still means redemptions route to the queue, but the crossing no longer self-amplifies.

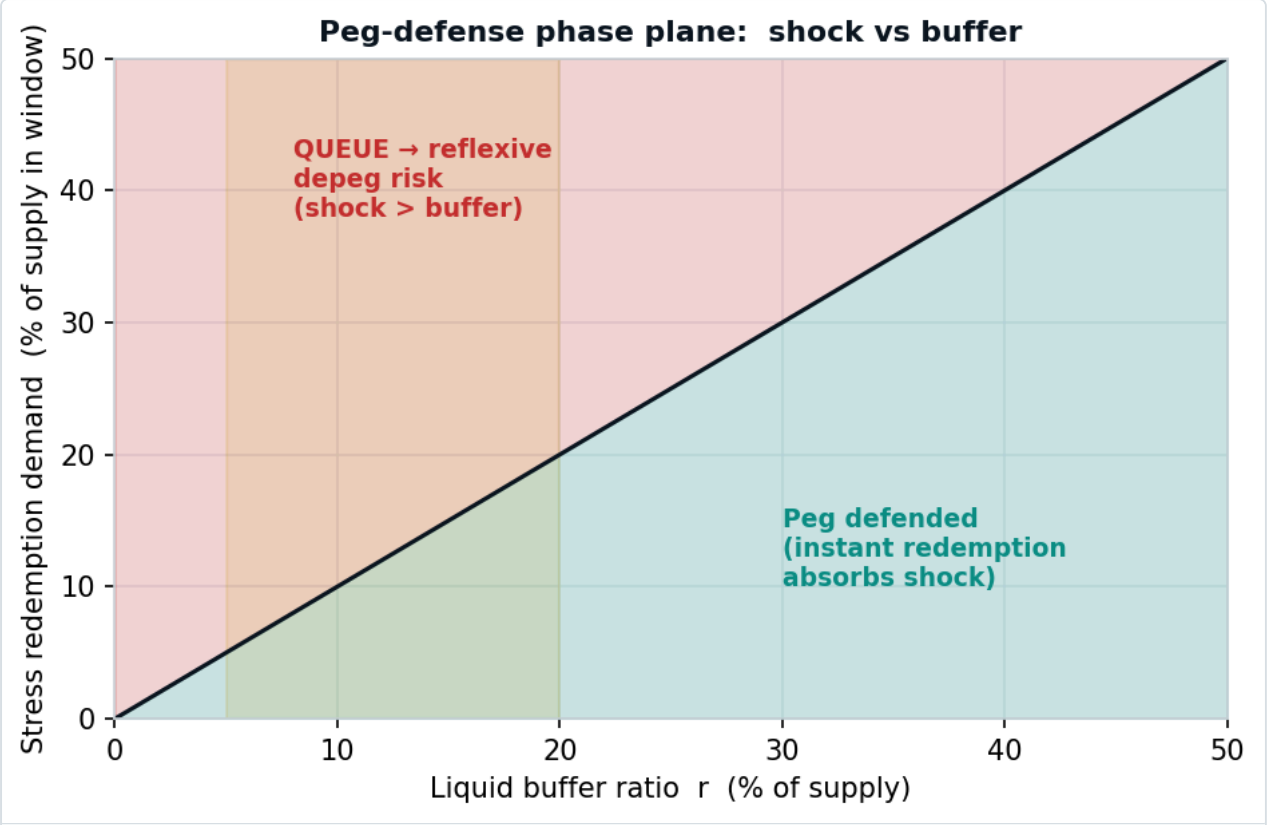

Fig. 3 — Peg-defense phase plane (liquidity boundary). The boundary is the line demand = r. Below it (teal) the instant buffer absorbs the shock and par arbitrage pins the peg. Above it (red) redemption routes to the FIFO queue and the secondary price reflects a *time-to-redemption* discount — but with a legible queue that discount is bounded and orderly, not a self-reinforcing spiral. The operator's job is to keep the operating point below the boundary by running r high enough that plausible shocks stay in the teal region.

> **Residual, precisely bounded.** The FIFO replay is a strong mitigation, not a guarantee, and two narrow items remain worth carrying: **(1)** ordering is *observed, not enforced* — completion still rests on the solver honoring the indexer order, with the improvement that any deviation is now detectable after the fact (a strong deterrent, but trust-but-verify); **(2)** the indexer becomes **availability-critical** — if it degrades during

> a crisis, exactly when queue load spikes, the legibility the peg now relies on degrades with it, so it warrants the same uptime/redundancy treatment as any consensus-critical service. Recommended hardening: publish on-chain buffer and aggregate queue-depth views so the guarantee does not depend solely on indexer liveness, run the indexer redundantly, and keep a path to on-chain FIFO/pro-rata enforcement as the eventual backstop. This vector is now **bounded (B-tier)**.

05 COLLATERAL-QUALITY DRIFT & THE FEE BRAKE

Because there is **no oracle in the mint/redeem path**, [redacted] treats every unit of USDC and USDT as exactly one dollar — predictable, but with a mechanical side effect: if one reserve asset depegs, the protocol becomes an arbitrage engine that trades good collateral for bad. An arbitrageur mints [redacted] by depositing the cheap asset (say USDT at $0.97) valued at par, then redeems for the sound asset (USDC at $1.00), also at par. Each cycle drains a unit of good collateral and leaves a unit of impaired collateral behind — Gresham's law, on-chain.

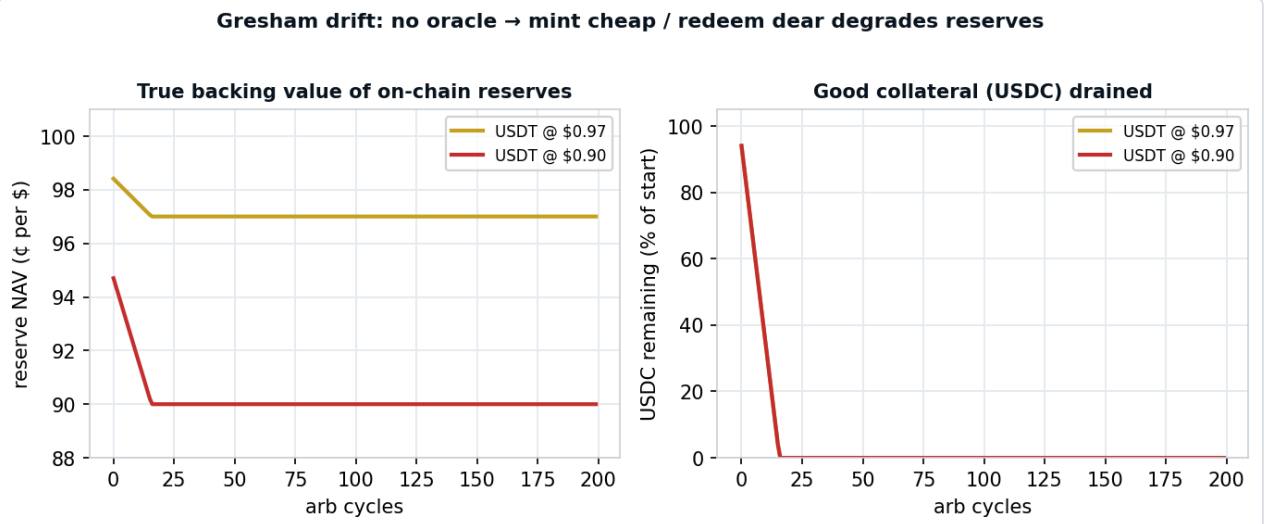

Fig. 4 — No-oracle collateral drift (fees off). Starting from a 50/50 USDC/USDT reserve, the arbitrage loop mints with the depegged asset and redeems the sound one. Left: true backing value of on-chain reserves (¢ per $) decays as good collateral is replaced by impaired collateral — mildly at a $0.97 depeg, sharply at $0.90. Right: the sound-asset (USDC) balance is drained toward zero within tens of cycles, while `collateralBalance` reports "fully backed" throughout because it counts units, not value.

Enter fees. The round-trip arbitrage profits only when the depeg gap exceeds the total mint + redeem fee: profit per dollar ≈ `(depeg gap) − (mint fee + redeem fee)`. A live fee therefore **hard-kills the arbitrage for any depeg smaller than the round-trip fee** — which covers essentially all routine USDC/USDT micro-wobble. The arbitrage only reappears for a large, genuine depeg, which is exactly the regime the opt-in circuit breaker (armed feeds, finding H-03) is for. Fees and the breaker **partition the problem cleanly**: fees handle the common small depegs, the breaker handles the rare large ones.

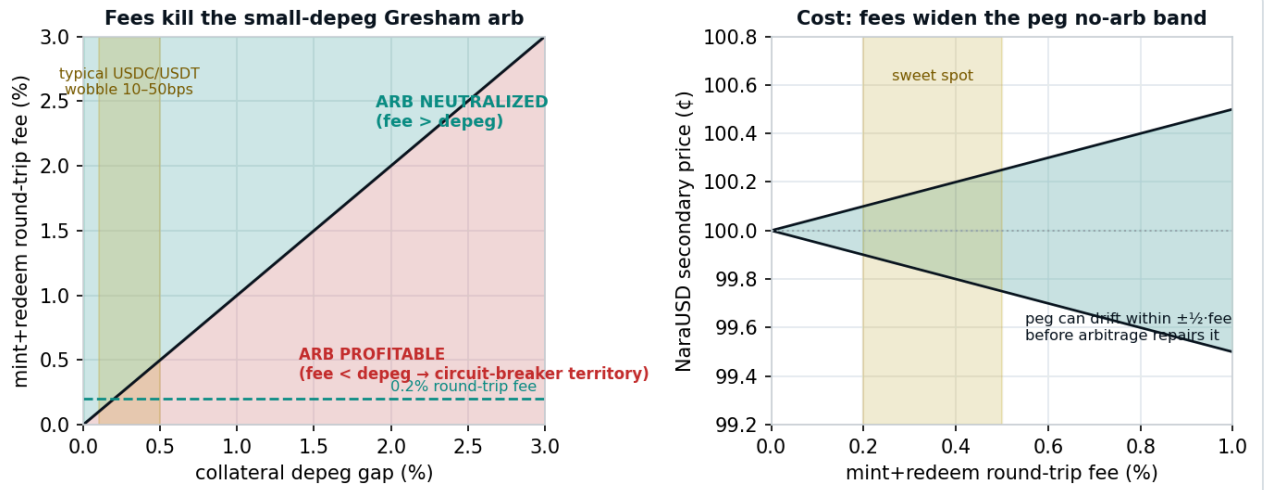

Fig. 5 — The fee brake and its cost. Left: the arbitrage is neutralized wherever round-trip fee > depeg gap (teal). A 0.2% round-trip fee closes any depeg up to 20bps — the whole typical stablecoin-wobble band — leaving only large depegs (circuit-breaker territory). Right: the cost. The same fee that kills the arbitrage also loosens the par-arbitrage that pins the peg, so [redacted] can drift within roughly ±½·fee before redemption arbitrage repairs it. The sweet spot is a round-trip fee of ~0.2–0.5%: tight enough to hold the peg inside ~±25bps, wide enough to neutralize routine depeg noise.

Two things to tune, one to bound. Set the fee in the ~0.2–0.5% round-trip band (kills micro-arb without a loose peg); this is the economic complement to arming the feeds (H-03), not a substitute for it — a large depeg still needs the breaker. And note that turning fees on **arms a live governance lever**: the 0→10% admin-settable range (E-13) means an operator could raise redeem fees under stress as a soft capital control — defensible as an emergency brake, but it should be **timelocked and range-bounded** so it cannot be quietly weaponized, the same theme as the upgrade-timelock item (H-02).

06 JUNIOR-TRANCHE RUN & COOLDOWN ADEQUACY

[redacted]'s loss waterfall makes [redacted]+ the **junior tranche**: losses hit the insurance fund first, then [redacted]+ stakers, sparing senior [redacted]. Junior holders are compensated for that risk with the >10% yield. The design is coherent — but it interacts with the loss-recognition lag Δ. When a loss is known off-chain but not yet booked via `burnAssets`, the share price is stale-high, and a holder who exits during Δ escapes their share of the loss entirely. The loss does not disappear — it is **re-concentrated onto the stakers who did not exit in time**, as `L/(1−f)`.

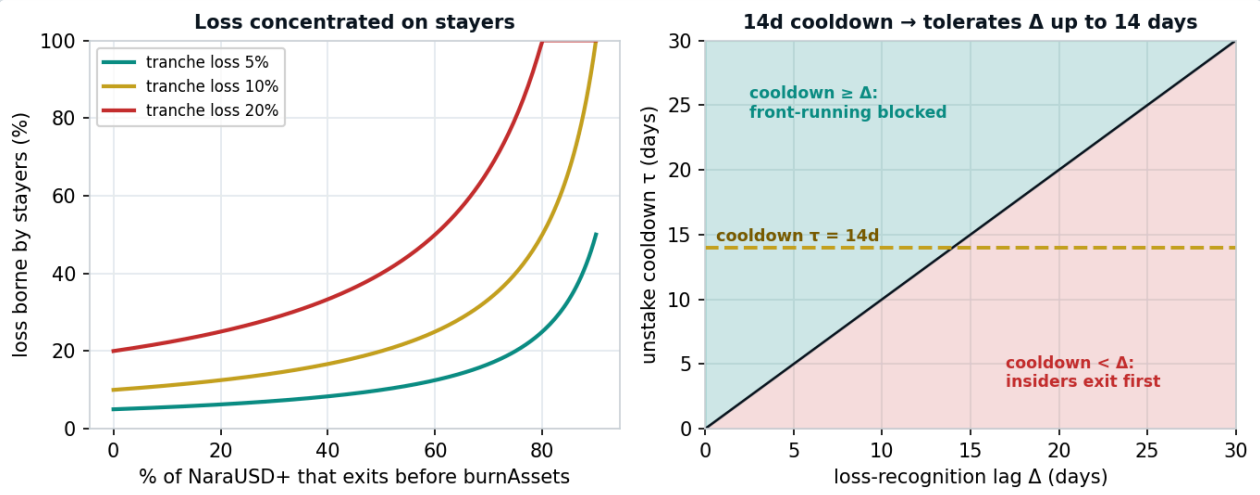

Fig. 6 — Loss concentration and the cooldown/lag race. Left: the loss borne by stayers is `L/(1−f)` where f is the fraction that exits first — a 10% tranche loss becomes 20% for stayers if half the tranche front-runs the write-down, and unbounded as f→1. Right: the safety condition is **cooldown τ ≥ lag Δ**. At τ = 14 days the safe (teal) region covers any loss-recognition lag up to a fortnight; front-running is blocked unless a known loss is left un-booked for more than two weeks.

The 14-day cooldown is the right fix for this vector. The concentration math (left) is unchanged — that is inherent to any junior tranche — but the *race* (right) now heavily favors the protocol. A staker can only front-run a loss if the recognition lag exceeds 14 days, a wide and realistic margin even for an impairment that takes time to confirm. To convert the margin into a guarantee, adopt an explicit policy — **book any impairment within a disclosed Δ_max well inside 14 days (e.g. 5–7)** — and disclose it. The auto-pause on `burnAssets` then closes the residual: it halts further exits the instant a loss is booked. Second-order benefit: a 14-day lock stickies junior capital, dampening the yield-chase reflexivity of E-04.

07 RANKED ECONOMIC ATTACK & FAILURE SURFACE

The complete surface at Rev-2 parameters, ranked by severity. **S-tier** can break the peg or cause an unrecoverable shortfall; **A-tier** cause material loss or a bounded depeg; **B-tier** are real but self-limiting. The parameter changes moved three vectors down a tier (E-02 → B, and E-05 / E-06 from existential to controlled-severe), leaving a **single S-tier item**. A recurring property persists: these vectors remain **coupled and multiplicative** — a residual Gresham drain (E-05) lowers effective backing, which lowers the shock threshold in the bank-run model (E-01) — so the lone existential vector is also the hub the others feed.

E-01 Liquidity-driven bank run (illiquidity without insolvency) — S-TIER

MECHANISM	Redemption demand in a 7-day window exceeds $r \cdot S$ + maturing credit; the solvent-but-illiquid protocol cannot honor instant redemption and routes to the queue. Untouched by any Rev-2 change.
DRIVER	Low buffer ratio r (yield pressure). See Figs. 1–2.
MITIGATION	**Publish and enforce a minimum-r floor** sized to a crisis-regime shock; check it as an invariant at `withdrawCollateral`; maintain a committed liquidity line; stagger PayFi maturities for a guaranteed daily inflow. The standout remaining action.

E-05 Gresham collateral-quality drain — large-depeg residual — A-TIER

MECHANISM	Mint-with-cheap / redeem-for-dear arbitrage drains the sound reserve asset while backing is reported at full units. See Figs. 4–5.
REV-2 STATUS	Small depegs (< round-trip fee) now **neutralized by fees**. Residual is the large-depeg case, still open until feeds are armed.
MITIGATION	Arm the opt-in per-collateral feeds + depeg circuit-breaker (H-03) for the large-depeg tail; keep the fee in the 0.2–0.5% band for the small-depeg body.

E-06 Junior-tranche front-running run — cooldown now 14d — A-TIER

MECHANISM	Informed holders unstake during the loss-recognition lag Δ, escaping their share; loss re-concentrates on stayers as $L/(1-f)$. See Fig. 6.
REV-2 STATUS	**Controlled** by the 14-day cooldown (safe while $\Delta < 14$d), pending an explicit Δ_{max} booking policy.
MITIGATION	Policy: book impairments within a disclosed Δ_{max} well inside 14 days; consider snapshotting NAV at loss-knowability rather than loss-booking.

E-03 Cross-chain peg fragmentation (hub/spoke liquidity split) A-TIER

MECHANISM	Redemption is hub-only; BSC-side [redacted] must bridge back to Ethereum to redeem. Under congestion or a bridge slowdown the spoke price can decouple from the hub price — two prices for "one asset".
DRIVER	Redemption centralized on the hub; bridge latency; thin spoke AMM depth.
MITIGATION	Seed and monitor spoke liquidity; publish hub-redemption-only clearly; consider spoke-side redemption intents.

E-04 Yield-chase reflexivity (TVL sensitivity to APR) A-TIER

MECHANISM	>10% APR attracts mercenary capital that is first to leave when yield dips or a competitor pays more; TVL that arrived for yield exits on yield, coinciding with the moments buffer is most needed.
REV-2 STATUS	Partly dampened — the 14-day cooldown stickies junior capital, so mercenary [redacted]+ cannot flee instantly.
MITIGATION	Size `r` to sticky-capital fraction, not total TVL; smooth yield; disclose reward discretion.

E-07 Staking-parameter risk (vesting / cooldown / fee levers) A-TIER

MECHANISM	6.5d vesting, 14d cooldown, and mint/redeem fees are all admin-settable (fees to a 10% cap; vesting/cooldown to 90d). A mis-set or hostile change alters run dynamics or enables timing games.
DRIVER	No timelock on parameter changes (ties to H-02); wider admissible ranges than the operating set.
MITIGATION	Timelock parameter changes; bound admissible ranges tighter than 0–90d and 0–10%; alert on change. Keep `vestingPeriod > 0` (I-05).

E-08 Uncapped collateral egress as a solvency lever A-TIER

MECHANISM	`withdrawCollateral` moves reserves off-chain to fund PayFi credit with no on-chain cap or timelock; the economic effect is to set `r` arbitrarily low in one transaction. Compromise or error directly de-funds redemption.
DRIVER	Governance lever; trust in off-chain accounting. Directly controls the E-01 master variable.
MITIGATION	Per-window egress cap; timelock large withdrawals; on-chain minimum-buffer invariant (`r_min`) checked at withdrawal.

E-09 Insurance-fund adequacy & opacity — A-TIER

MECHANISM	The first loss-absorbing layer is an off-chain insurance fund of undisclosed size. If it is small relative to the PayFi book, junior stakers are effectively first-loss and don't know it.
DRIVER	Off-chain, unattested fund size.
MITIGATION	Disclose and attest insurance-fund size and its ratio to credit exposure; set a minimum coverage target. Fee revenue is a natural funding source.

E-10 Concentration / single-obligor PayFi default — A-TIER

MECHANISM	PayFi credit concentrated in few counterparties → a single default is a large, discrete loss that can exceed the insurance fund and hit [redacted]+ directly.
DRIVER	Off-chain credit-book concentration (unobservable on-chain).
MITIGATION	Disclose obligor concentration limits; diversify; publish a credit-policy summary.

ID	BOUNDED (B-TIER) VECTOR	WHY SELF-LIMITING
E-02	Redemption-queue depeg **(downgraded from S-tier)**	Indexer-enforced FIFO removes the informational asymmetry / panic amplifier. Residual = solver-follows-index trust + indexer liveness (§04).
E-11	Per-block mint/redeem rate-limit gaming (1M/block)	Slows but does not prevent a large drain; bounded by block cadence.
E-12	Vesting-window donation / share-price timing	6.5-day vesting + `totalAssets` minus unvested neutralizes the sandwich; longer vesting only strengthens this.
E-13	Fee-lever weaponization (0% → up to 10%)	Now that fees are live this is a real lever, but capped at 1000bps and visible; bound + timelock the range.
E-14	Secondary-market thin-liquidity depeg (non-redemption)	Par arbitrage repairs it within the fee band as long as instant redemption is open; bounded by `r`.
E-15	Cross-composer refund griefing	Compose path is binary (complete or refund); loss bounded to gas + LZ fees.

08 QUANTITATIVE MODELS — REFERENCE

The seven models underpinning this addendum, with their governing equations and the parameter that dominates each. All are reproducible from the accompanying `sims.py` / `sims2.py`.

ID	MODEL	GOVERNING RELATION	DOMINANT PARAM
M1	Liquidity coverage	`cover = r·S; days = r·S / (d·S − (1−r)S/D)`	`r, D`
M2	Bank-run probability	`P(Σ₇ demand > r·S + 7·inflow)`	`r, shock p`
M3	Reflexive/liquidity phase plane	`boundary: demand = r`	`r`
M4	Gresham drift	`NAV = (u_c·1 + u_t·p)/(u_c+u_t)`	`depeg p`
M5	Junior loss concentration	`stayer loss = L/(1−f)`	`f, Δ`
M6	Cooldown adequacy	`safe ⇔ τ ≥ Δ  (τ = 14d)`	`Δ, τ`
M7	Fee arbitrage / peg-band	`arb dead ⇔ f_m+f_r > depeg; band = ±½(f_m+f_r)`	`fee, depeg`

> **Model scope & honesty.** These are structural mechanism models, not forecasts. They take deliberately conservative but plausible parameter ranges and show *where* the risk lives and *which knob* moves it, so the client can set policy on the knobs. They do not claim probabilities of real-world events; the buffer ratio `r`, PayFi maturity ladder, insurance-fund size, loss-recognition timing, and the exact fee level are off-chain or configurable quantities we swept across ranges. Supplying the live figures would let us tighten every model to the current book.

09 ECONOMIC HARDENING — PRIORITY ORDER

Re-ranked for Rev-2. With the cooldown, fees, and queue-legibility items now addressed, the buffer floor rises to the top and the fee tuning drops to a settings task.

#	ACTION	CLOSES	COST
1	**Publish & enforce a minimum buffer ratio r_min** sized to a crisis-regime shock; check it as an on-chain invariant at `withdrawCollateral`. The lone remaining existential item.	E-01, E-08	Med
2	**Arm per-collateral price feeds + depeg circuit-breaker** for the large-depeg tail that fees do not cover.	E-05 / H-03	Low
3	**Set the mint/redeem fee in the ~0.2–0.5% sweet spot;** timelock and range-bound the 0–10% fee lever so it can't be weaponized.	E-05, E-13	Low
4	**Adopt & disclose Δ_max well inside the 14-day cooldown** (e.g. book impairments within 5–7 days).	E-06	Low
5	**Back the FIFO queue with on-chain buffer & queue-depth views + a redundant indexer;** keep a path to on-chain FIFO/pro-rata enforcement.	E-02	Med
6	**Per-window collateral-egress cap + timelock** (complements the r_min invariant).	E-08 / H-01,H-02	Med
7	**Disclose & attest insurance-fund size and PayFi obligor concentration;** route fee revenue to the fund.	E-09, E-10	Low
8	**Seed & monitor spoke (BSC) liquidity;** document hub-only redemption.	E-03	Low

Items 2, 3, 4, 7 and 8 are low-cost settings and disclosures achievable immediately. Items 1, 5 and 6 are the structural changes that move [redacted] from "solvent but runnable" to "solvent and run-resistant." After the Rev-2 hardening, **item 1 — the buffer floor — is the single thing standing between the current design and genuine run-resistance.**