A Delta-Neutral Meme Treasury

Formal guarantees, and their limits, for a fee-financed hedged market maker on Robinhood Chain

Hedgehood · 5 September 2026 · working paper, v1
Reproducible arithmetic: `python3 docs/paper/verify_numbers.py` · system of record: `docs/MEME_TREASURY.md`

Abstract

We analyse a treasury design in which a memecoin's creator tax, denominated in a tokenised equity (NVDA on Robinhood Chain), is routed without human signature into two legs: perpetual-futures margin on Lighter and concentrated-liquidity inventory in the NVDA/USDG pool, with the resulting spot exposure continuously shorted. We ask what part of this is *provable*. Four properties are: (i) delta neutrality is an exact algebraic identity, not an approximation, and it holds only if exposure is summed over all six locations treasury stock can occupy (Thm. 3–4); (ii) the liquidation distance of the hedge has a closed form that is *independent of the size of the fee flow*, so a solvency margin chosen once remains valid at any scale — at the deployed 40% it absorbs a $+62\%$ overnight gap in the worst case (Thm. 8, Cor. 9–10); (iii) the keeper's reachable set is confined to three addresses all owned by the Safe, and the value a fully compromised keeper can leak per day is bounded by $2C_m\theta$ — \$168/day at the deployed parameters (Thm. 12–14); (iv) a buyback financed strictly out of realised yield cannot deplete principal, and buy-and-burn strictly raises the pool's total-capitulation floor price (Thm. 16–17). Profitability, by contrast, is *not* a theorem: it reduces to one measurable inequality, $\varphi\dot{V} > \sigma^2 x_q/4$ — fee revenue against the pool's own loss-versus-rebalancing — which the NVDA/USDG 0.05% pool currently passes by a factor of 3.0 today — it read 2.0 thirteen hours earlier, which is the point — with a closed-form profit-maximising liquidity share $s^\star = 1 - \sqrt{\ell/\varphi\dot{V}} = 42\%$ (Thm. 5, Criterion 7). We give the pairing-eligibility criterion that follows from reflexivity (Criterion 19), which quantitatively justifies pairing against NVDA (\$537k of flow to reach the gate) rather than a thin token (\$785), and we close with the properties the design does *not* have: it does not manufacture turnover, and one event — a global pause of the equity token — breaks the hedge's liquidity, though not its bound.

摘要

本文分析的对象：一个 meme 的创作者税以股票代币（Robinhood Chain 上的 NVDA）结算，无需人工签名地被拆成两条腿 —— Lighter 上的永续保证金与 NVDA/USDG 池里的集中流动性库存 —— 并把由此产生的现货敞口持续做空。我们问的是：这里面哪些部分是可以*证明*的。可证的有四条：（一）delta 中性是一个精确的代数恒等式而非近似，且只有把国库股票可能停留的全部六个位置都计入敞口时才成立（定理 3–4）；（二）对冲的清算距离有闭式解，且*与费用流的规模无关*，所以一次选定的保证金比例在任何规模下依然有效 —— 按已部署的 40%，最坏情况下能吃住 $+62\%$ 的隔夜跳空（定理 8、推论 9–10；注意这个跳空幅度是我们自己选的对冲目标，不是交易所的保证）；（三）keeper 的可达集被限制在三个地址，且都由 Safe 拥有，一个完全被攻破的 keeper 每天最多泄漏 $2C\theta$，按已部署参数每天不超过 420 美元、实际上界 168 美元（定理 12–14）；（四）只用已实现收益做回购不可能吃掉本金，且回购销毁严格抬高池子的"全员抛售地板价"（定理 16–17）。相反，*盈利不是定理*：它归结为一个可测的不等式 $\varphi\dot{V} > \sigma^2 x_q/4$（手续费收入对池子自身的 loss-versus-rebalancing），NVDA/USDG 0.05% 池今天以 3.0 倍通过 —— 十三小时前还是 2.0 倍，这恰恰是重点 —— 并有闭式最优份额 $s^\star = 1 - \sqrt{\ell/\varphi\dot{V}} = 42\%$（定理 5、判据 7）。最后给出由反身性导出的配对资产合格判据（判据 19）：把池价推到 50bp 门槛需要 53.7 万美元的净流入（NVDA）对 785 美元（薄池），这从数量上支持了选 NVDA 而拒绝薄池标的。以及本设计*不具备*的性质：它不制造交易量；并且有一个事件 —— 股票代币全局暂停 —— 会破坏对冲的流动性（虽然不破坏它的边界）。

1. What "correct" is being claimed

On-chain memecoin treasuries are, almost without exception, wallets holding the asset they were paired against: one prominent NVDA-paired launch holds 762 NVDA and nothing else. Such a treasury is a levered directional bet on the pairing asset wearing the clothes of a treasury. The design analysed here differs in exactly one structural respect:

the inventory is *hedged*, and the hedge is maintained by a keeper whose authority is bounded by immutable contract parameters rather than by trust.

That claim invites a sharper question than "does it work". A design of this shape makes four separable promises, and they have very different epistemic status:

1. **Neutrality.** The treasury's value does not move with the price of the pairing asset. *This is provable, exactly, and we prove it.*
2. **Solvency.** The hedge survives any adverse move smaller than a stated bound. *Provable given the venue's maintenance fraction; we give the bound, show it is scale-free, and are explicit that the gap it is sized against is a choice rather than a venue guarantee.*
3. **Custody.** Automation does not create a path for value to leave. *Provable as a reachability property of the deployed contracts, with a quantitative leakage bound.*
4. **Yield.** The hedged position earns more than it costs. *Not provable. It is a measurable inequality about a specific pool, which can and does change.*

Conflating the fourth with the first three is the characteristic error of this genre. We keep them apart throughout: results labelled `Theorem` are identities or bounds that follow from the contracts and the venue's rules; results labelled `Criterion` are empirical inequalities that must be re-measured. Section 9 gives the current measurements; Section 10 lists what remains unproved.

2. The system

Three venues and one owner. A 2-of-3 Safe $\mathcal{S}$ owns everything. A memecoin launched on Pons v2 is quoted in a tokenised equity X (NVDA); its creator tax accrues in X to a pull-based escrow, is claimed by a permissionless `FeeSink` that can only forward to $\mathcal{S}$, and is then split by `HedgeMarginRouter` into perpetual margin on Lighter's Robinhood Chain instance (an account owned by $\mathcal{S}$) and inventory in a `SafeLpController` that holds a Uniswap v3 NVDA/USDG position. An off-chain keeper, holding a trade-only API key and an EOA that is the router's `keeper`, shorts the resulting exposure on the perpetual.

Symbol	Meaning	**Value** (on chain unless marked)
P	fair price of X in USDG (oracle / perp mark)	`231.18` *measured*
φ	pool fee tier	`5 bp`
$L, [P_a, P_b]$	our v3 liquidity and range	controller, `recenter`
$x_q \equiv L\sqrt{P}$	virtual quote reserve ("depth")	`see §9`
E	treasury exposure to X, in shares	`rhsvc/treasury.py`
q	perp position, signed (short < 0)	Lighter acct. 22669, *flat*
m	`marginBps`: share of each allocation sold for margin	`0.40`
μ	venue maintenance-margin fraction	`0.03`
G	closed-market gap the hedge is sized against *(assumed, not a venue guarantee — see Cor. 10)*	`0.158`
C	`dailyCapUsdg`, per fixed 24 h window	`20,000`
θ	`maxDeviation + maxSlippage + poolFee/100`	`105 bp`
φ_c	Pons curve fee	`100 bp`
τ	Pons creator tax	`333 bp` *planned*
q_0, $X^\star$	curve phantom reserve, graduation threshold	`16.64, 41.6 NVDA`

Table 1. Notation. Every value is read from chain at the time of writing except the two marked: P is a measurement, and τ is *planned, not deployed* — no token has been launched (`config/hedge.json`, `launch_token: null`), so nothing has ever paid this tax. The perpetual account

exists and is flat; the controller holds no position. Nothing in this paper describes money at risk today. Router `0x9350…D536` (v2), NVDA controller `0x0F7e…4d7b`, Safe `0xB5C9…489b`. The cap was raised from 2,000 to 20,000 USDG on 2026-09-05 under a router whose runtime code differs from its predecessor's in exactly the four occurrences of that immutable; Lemma 13 and Theorem 14 make the leakage bound linear in that cap, while Corollary 9 is what leaves the solvency bound of Theorem 8 untouched.

Six locations can hold treasury stock at any instant, and the distinction matters in Theorem 2: the escrow (unclaimed), the `FeeSink` (claimed, unforwarded), the Safe (forwarded, unallocated), the controller's idle balance, the v3 position's principal, and the position's uncollected fees. Write these e, f, s, i, g and $x(P)$ respectively, all in shares, so that

$$E \;=\; x(P) + i + g + s + f + e.$$

3. Revenue is a function of turnover, not of price

Pons v2 takes the creator tax off the quote leg of every trade: a buy of gross quote size a contributes τa, and a sell whose gross quote proceeds are b contributes τb. Nothing about the token's price path enters.

> **Theorem 1 (turnover invariance).** Let $\mathcal{T}_t$ be gross quote-denominated turnover on the curve and its post-graduation hook up to time t. Then treasury revenue is $R_t = \tau \mathcal{T}_t$, a non-decreasing process independent of the token's price path, of its price level, and of the order of trades.
>
> *Proof.* Revenue is a sum of per-trade terms each equal to τ times that trade's gross quote leg, with no dependence on the reserve state; summing over any permutation of the same trade set gives the same total, and each term is non-negative. □

Two consequences are worth naming because they are frequently assumed away. First, the treasury is *long volume and flat direction*: a token that round-trips to zero having traded $10M pays the same $333k as one that goes to $10M on the same turnover. Second, revenue is exogenous: the only quantity in Theorem 1 that the treasury does not control is the one that determines its size. No mechanism in this paper manufactures $\mathcal{T}$.

> **Corollary 2 (graduation).** With a constant-product curve seeded by a phantom quote reserve q_0 and a graduation threshold $X^\star$ on the accumulated *net* quote reserve, the fraction of supply sold on the curve is $X^\star/(q_0 + X^\star)$, independent of the path. The tax collected on the way is at least $\tau X^\star/(1 - \varphi_c - \tau)$, with equality on a monotone-buy path.
>
> *Proof.* Constant product on $(q_0 + x)(T_0 - y) = q_0 T_0$ gives $y = T_0\, x/(q_0 + x)$, a state function of the reserve; at $x = X^\star$ the sold fraction is $X^\star/(q_0 + X^\star)$, whatever order the trades arrived in. The tax is different: each unit of gross inflow contributes $1 - \varphi_c - \tau$ to the reserve, so a purely monotone ascent needs gross $X^\star/(1 - \varphi_c - \tau)$ — but a sell pushes the reserve back down, and by Theorem 1 tax accrues on gross turnover in both directions, so any path with sells pays strictly more to reach the same threshold. The formula is a floor, not an identity. □

Both constants are read from the chain, not assumed. `pairTokenEconomics(NVDA)` on the Pons factory `0x7eD5…EC7e` returns $(16.64 \times 10^{18},\ 41.6 \times 10^{18},\ 18)$ — phantom reserve, threshold and decimals. Their ratio is exactly $5/2$, which is what makes the sold fraction exactly $5/7$. The threshold is read here as applying to the *net* reserve; on a gross reading the figures would be 70.5% and $319 instead.

At the deployed parameters: 71.43% of supply sells on the curve, gross inflow to graduation is 43.48 NVDA, and the creator tax collected is 1.448 NVDA — $334. This is the correct order of magnitude to hold in mind: the launch itself funds almost nothing. Every $1M of subsequent turnover pays $33,300. The treasury's fortune is made or not made after graduation, which is why the yield machinery below is a secondary term, not the thesis.

4. Neutrality is exact, and requires all six locations

Let $W(P)$ be the treasury's value in USDG. Inside the range, the v3 position holds $x(P) = L(P^{-1/2} - P_b^{-1/2})$ shares and $y(P) = L(\sqrt{P} - \sqrt{P_a})$ quote, so its value is $V(P) = y + xP = L\left(2\sqrt{P} - \sqrt{P_a} - P/\sqrt{P_b}\right)$. Adding the linear legs,

$$W(P) \;=\; V(P) \;+\; \big(i+g+s+f+e\big)P \;+\; M \;+\; q\,(P-P_0) \;+\; (\text{USDG holdings}).$$

Theorem 3 (exact neutrality to a common move). For $P \in [P_a, P_b]$, $\partial W/\partial P = E + q$. Hence the hedge $q^\star = -E$ makes the treasury's value stationary in the price of X exactly, not to first order in a small-move approximation.

Proof. Differentiate: $\partial V/\partial P = L\big(\tfrac{1}{2}\cdot 2P^{-1/2} - P_b^{-1/2}\big) = L(P^{-1/2} - P_b^{-1/2}) = x(P)$. (The restriction to $[P_a, P_b]$ is conservative rather than necessary: above P_b the position is all quote, $x = 0$ and V is constant; below P_a it is all stock, $V = xP$ with x constant. $V' = x$ in both, so the identity holds on the whole line.) The remaining stock legs are linear with slope equal to their quantity, the perp mark-to-market has slope q, and margin and USDG have slope zero. Summing, $\partial W/\partial P = x(P) + i + g + s + f + e + q = E + q$, which vanishes iff $q = -E$. □

The identity $\partial V/\partial P = x(P)$ — the position's delta *is* its stock balance — is what makes the hedge implementable by a keeper that can only read balances. It needs no volatility model, no Greeks, and no oracle beyond the one used to price the balance.

One symbol, three prices. W above is written in a single P, and the three legs are not marked at the same one: the v3 position is valued at the *pool* price, the bare stock legs at the *oracle*, and the perpetual's mark-to-market at the venue's *mark*. Honestly,

$$dW \;=\; x\,dP_{\text{pool}} \;+\; (i+g+s+f+e)\,dP_{\text{oracle}} \;+\; q\,dP_{\text{mark}},$$

so $q = -E$ annihilates it only for a *common* move of the three. The residual is $x\,(dP_{\text{pool}} - dP_{\text{mark}}) + (i+g+s+f+e)(dP_{\text{oracle}} - dP_{\text{mark}})$ — basis, not direction. This paper cannot wave that away, because it monetises the same basis elsewhere: the deviation gate exists because pool and oracle diverge (§4.1), Proposition 18 is a trade on $\bar{P}_{\text{pool}} \neq P$, and §9's disagreement between the two LVR estimates is diagnosed as the mark not tracking the pool. So Theorem 3 should be read as what it is: the hedge removes *direction* exactly and leaves basis risk untouched. That residual is listed in §10 with the others.

Theorem 4 (necessity of the full exposure sum). If the keeper hedges $q = -x(P)$, i.e. the LP leg only, the residual delta is exactly $i+g+s+f+e > 0$. Moreover the residual is *maximised precisely on the days the router's deviation gate is active*.

Proof. The first claim is Theorem 3 with $q = -x(P)$. For the second: `allocate` reverts when $|P_{\text{pool}} - P|/P > \delta_{\text{gate}}$, so during any interval in which the pool is off the oracle, arriving fees accumulate in s (and in f, e) and cannot be moved into i or x. Premium episodes are exactly the episodes in which the gate binds, so $\sup_t(s+f+e)$ is attained inside them. □

This is not a hypothetical. It is the reason `rhsvc/treasury.py` sums six components at one pinned block and raises rather than defaulting a component to zero: a silently dropped component is an under-hedge that trips no alarm, and it would be dropped on the day the position is largest and the pool is most dislocated. The fail-closed read is a correctness requirement, not defensive style.

The second half of Theorem 4 is contingent on a symmetric gate, and that is being changed. The deployed router (v2) refuses to sell whenever $|P_{\text{pool}} - P|/P > \delta_{\text{gate}}$ in either direction, which is why premium episodes and gated episodes coincide. An open change (router v3) makes the gate *one-sided*: only a pool below the oracle blocks a sale, selling into a premium is unbounded. The motivation is exactly the weekend case — Chainlink is frozen, the pool carries the premium, and a symmetric gate refuses to sell into the very dislocation the inventory is held for. Under v3 the first half of Theorem 4 (the residual is $i+g+s+f+e$) is untouched, but the second half flips: the residual is then maximised on *discount* days, not premium days. The keeper must still sum all six either way, which is the claim that matters; only the identification of the worst day moves.

4.1 What neutrality does not buy: the gamma rent

Delta neutrality removes the *variance* of the price term. It does not remove its *drift*. The v3 value function is strictly concave, $\partial^2 V/\partial P^2 = -L/(2P^{3/2}) < 0$, so a continuously re-hedged position bleeds at the Itô rate $\tfrac{1}{2}V''(P)\sigma^2P^2$.

Theorem 5 (gamma rent / LVR). Under $dP = \sigma P\,dW$ and continuous re-hedging, the expected loss rate of the hedged position from price dynamics alone is

$$\ell \;=\; \tfrac{1}{4}\,\sigma^2 L\sqrt{P} \;=\; \tfrac{1}{4}\,\sigma^2 x_q,$$

independent of the range endpoints while the price is inside the range.

Proof. $\tfrac{1}{2}V''\sigma^2P^2 = \tfrac{1}{2}\big(-L/(2P^{3/2})\big)\sigma^2P^2 = -\tfrac{1}{4}\sigma^2L\sqrt{P}$. Neither P_a nor P_b appears; they enter only through L for a given capital, and through the boundary at which the formula stops applying. □

Consistency check. A full-range (v2) position has value $2L\sqrt{P} = 2x_q$, so Theorem 5 reads $\ell = \sigma^2 V/8$ — the standard loss-versus-rebalancing rate of Milionis, Moallemi, Roughgarden and Zhang. Theorem 5 is that result specialised to a concentrated range and expressed in the one quantity a keeper can read on chain.

Do not double-count. Our calculator (`lp_calc.py`) subtracts a *measured* markout — realised execution price against the perp mark five minutes later, per unit of volume — and carries no separate LVR term. That is correct, not an omission: under a martingale mark and tight hedging, markout-against-mark is a realised-path estimator of the same quantity Theorem 5 gives in expectation. Adding both would charge the position twice. The two estimates should be reported side by side precisely because they can disagree (§9), and their disagreement is information about the mark, not about the position.

4.2 Concentration is not alpha; there is a closed-form optimal share

Theorem 6 (scale invariance of the edge in concentration). While our liquidity share $s = L/(L_p + L)$ is small, both fee income $\varphi\dot{V}s$ and gamma rent $\tfrac{1}{4}\sigma^2L\sqrt{P}$ are linear in L. Hence tightening the range at fixed capital multiplies income and cost by the same factor, and cannot change the sign of the edge.

Proof. $s \approx L/L_p$ for $L \ll L_p$, so fees $\propto L$; Theorem 5 gives cost $\propto L$. The ratio is invariant. □

Concentration therefore buys magnitude and pays for it in rebalancing frequency and truncation risk — it is not an edge. What determines the sign is the pool, not our range. Dropping the small-share approximation gives the sharper statement, and it is the practically useful one:

Criterion 7 (entry, and optimal share). Let $\Pi(L) = \varphi\dot{V}\,\frac{L}{L_p+L} - \tfrac{1}{4}\sigma^2\sqrt{P}\,L$. Then Π has an interior maximum at $L^\star = \sqrt{4\varphi\dot{V}L_p/(\sigma^2\sqrt{P})} - L_p$, which is positive if and only if

$$\varphi\dot{V} \;>\; \tfrac{1}{4}\sigma^2 x_q^{\text{pool}},$$

i.e. iff the pool as a whole earns more in fees than it loses to rebalancing. The optimal share is

$$s^\star \;=\; 1 - \sqrt{\ell_{\text{pool}}/\varphi\dot{V}}\,.$$

Proof. $\Pi'(L) = \varphi\dot{V}L_p/(L_p+L)^2 - \tfrac{1}{4}\sigma^2\sqrt{P} = 0$ gives $(L_p+L)^2 = 4\varphi\dot{V}L_p/(\sigma^2\sqrt{P})$, hence $L^\star$; $L^\star > 0$ iff $4\varphi\dot{V} > \sigma^2\sqrt{P}L_p$, which is the displayed inequality since $x_q^{\text{pool}} = L_p\sqrt{P}$. Substituting $L^\star$ into $s = L/(L_p+L)$ gives $s^\star$. □

What $s^\star$ assumes. The optimisation holds $\dot{V}$ and σ fixed while we add L. That is harmless at our current 0.3% share and heroic at $s^\star$: reaching 42% means adding about 72% to the pool's in-range liquidity, which would tighten the spread arbitrage pays and plausibly move both volume and the pool's own price process. Read $s^\star$ as an upper bound on how far it is worth growing, not as a target that survives being reached.

Three things follow that are worth stating plainly. The entry test does not involve our own size: a pool is either worth making or it is not, and if it is not, no amount of cleverness about ranges fixes it. The APR on capital is maximised as capital $\to 0$ (the first dollar earns the undiluted fee share), while total profit is maximised at $s^\star$; a treasury that starts small and grows out of fees therefore lives in the high-APR / low-total regime for a long time, which is the correct regime to be in. And $s^\star$ depends on the fee-to-LVR ratio only, so it can be recomputed from public pool data every day.

4.3 Discrete hedging

The keeper re-hedges when $|q + E| > b\,|E|$ for a band b, subject to a venue minimum order size. Residual delta is bounded by $b|E|$ by construction, so the tracking error over a horizon T is $O(bE\sigma P\sqrt{T})$; the classical band-hedging

trade-off makes the optimal b increase with trading cost. On Lighter's Robinhood Chain NVDA market, taker fees are zero, so the cost term vanishes and the optimal band collapses to the venue's granularity floor (0.04 shares / 10 USDG). The keeper's band is therefore not a risk–cost optimum but a discretisation artefact, which is the right answer when execution is free.

5. Solvency: the liquidation bound is scale-free

Each allocation sells a fixed fraction m of the arriving stock for USDG margin and sends $1-m$ to the LP controller, *before* any price is observed on the LP leg. Let a cohort of S shares arrive at price P, and let $\rho \in [0,1]$ be the fraction of the LP leg that is actually held as stock (the v3 range holds a mix; $\rho \approx \frac{1}{2}$ just after a symmetric re-centre, $\rho \to 1$ at the bottom of the range). The short notional is $N = (1-m)\rho SP$ and the margin is $M = mSP$.

> **Theorem 8 (liquidation distance).** A short of notional N at entry P_0 with margin M and maintenance fraction μ is liquidated at $P_{\text{liq}} = (P_0 + M/q)/(1+\mu)$, i.e. at a relative distance
>
> $$d \;=\; \frac{1 + M/N}{1+\mu} - 1 \;=\; \frac{M/N - \mu}{1+\mu} \;=\; \frac{1}{1+\mu}\left(\frac{m}{(1-m)\rho} - \mu\right).$$
>
> *Proof.* Equity at price P is $M + q(P_0 - P)$ for short size $q > 0$; maintenance requires $M + q(P_0 - P) \geq \mu q P$, i.e. $P \leq (P_0 + M/q)/(1+\mu)$. Dividing by P_0 and using $M/N = M/(qP_0)$ gives the displayed d. □

> **Corollary 9 (scale-freeness).** d depends on m, ρ, μ only. It is independent of S, hence of the size of the fee flow, of the token's success, and of time.
>
> *Proof.* $M/N = m/((1-m)\rho)$; S and P cancel. □

This is the single most useful property of the split. A margin ratio validated once at $300 of fees remains valid at $300,000, because the ratio is imposed on every cohort at arrival rather than reconciled after the fact. At the deployed $m = 0.40$, $\mu = 0.03$:

state	ρ	M/N	distance d
just after a symmetric re-centre	0.50	`1.333`	`+126.5%`
price at range bottom, LP entirely stock (worst case)	1.00	`0.667`	`+61.8%`

Table 2. Liquidation distances from Theorem 8. Both figures were independently derived in `docs/MEME_TREASURY.md` §2 as +126% and +62% before this paper existed; the agreement is a check on the formula, not a coincidence.

> **Corollary 10 (surviving a gap of G).** The position survives an adverse move of G over a window in which it cannot be re-margined provided $m \geq \dfrac{\rho A}{1+\rho A}$, where $A = (1+G)(1+\mu) - 1$. At $G = 0.158,\ \mu = 0.03,\ \rho = 1$ the requirement is $m \geq 16.16\%$; the deployed $m = 40\%$ carries $2.48\times$ that, and G is 0.26 of the worst-case distance.
>
> *Proof.* Require $1 + d \geq (1+G)$, i.e. $1 + m/((1-m)\rho) \geq (1+G)(1+\mu)$; solve for m. □

> **Where $G = 0.158$ comes from, and what it is not.** An earlier draft called this a cap the perpetual venue imposes. It is not, and the distinction matters. 15.8% is the far end of Hyperliquid's *discovery-bound* ratchet — the mark is held within $\pm 1/\lambda_{\max}$ of a reference price and re-anchors at 90% of that bound, so several ratchets reach roughly $\pm 15.8\%$ — and it governs the GME line on that venue, not the NVDA line analysed here. **Lighter is not known to cap the weekend mark at all**: its NVDA market publishes a maintenance fraction (μ: `maintenance_margin_fraction 300`, confirming $\mu = 0.03$) and funding clamps, and no mark bound. So Corollary 10 is a statement about a gap size we choose to size against, not a guarantee inherited from the venue. Read it as: *the hedge survives a $+62\%$ overnight gap in the worst LP composition, and $+126\%$ in the ordinary one.* That is the honest claim, and for a mega-cap it is a wide one — but it is a claim about how large a gap we tolerate, not about how large a gap is possible.

The chosen $2.5\times$ headroom is deliberate and its justification is an *absence*, not a preference: `topUpMargin` can only pull USDG that is already in the Safe, and USDG returns to the Safe only through a signed sweep. Until that gap is closed, the front-loaded 40% is the only automatic margin defence there is, so it is sized to survive a large overnight gap without help. Once an automatic USDG source exists, the bound above says any $m \geq 16.2\%$ survives a 15.8% gap, so it does not distinguish 30% from 25% at all. What does is the keeper's own top-up trigger, which fires at a liquidation distance of 20%: solving $d(m, \rho = 1) = 0.20$ puts that at $m \approx 19.1\%$, so $m = 25\%$ clears it ($d = 29.4\%$,

1.47× the trigger) and $m = 30\%$ clears it more comfortably ($d = 38.7\%$, 1.93×). An earlier draft said 25% "sits under the 20% trigger", which confused a bound on m with a bound on distance; the design doc's own wording is that 25% is *close to* the trigger, not under it.

> **The static worst case is conservative.** Reaching $\rho = 1$ requires the price to fall to the bottom of the range, and a fall credits the short with unrealised profit, raising M at the same time as it raises ρ. The $+61.8\%$ figure evaluates both at the same price and therefore understates the true distance; the honest worst case is the composition — fall to the range bottom, then gap up — which is what the design doc's +62% describes.

5.1 Two corrections to the keeper's margin rules

The keeper's tripwires (`rhsvc/guards.py`) use linearised versions of Theorem 8. Both are worth checking rather than assuming, and they behave differently:

> **Proposition 11.** (a) The exact margin needed to move the distance from d_c to d_t is $N(1+\mu)(d_t - d_c)$; the implemented rule $N(d_t - d_c)$ under-funds by the factor $1+\mu$ — 3.0% at $\mu = 0.03$. (b) The emergency-reduction rule "keep the fraction d_c/d_t of the size" is *conservative*: whenever it fires it reduces at least as much as the exact requirement $M/N/(d_t(1+\mu)+\mu)$.
>
> *Proof.* (a) From Theorem 8, $M = N(d(1+\mu)+\mu)$; differencing at fixed N gives $\Delta M = N(1+\mu)(d_t - d_c)$. (b) Writing $c = M/N$, the code keeps $k_g = d_c/d_t$ with $d_c = (c-\mu)/(1+\mu)$, while exactness requires $k_e = c/(d_t(1+\mu)+\mu)$. Then $k_g \le k_e$ reduces to $c \le d_t(1+\mu)+\mu$, which is exactly the condition $d_c \le d_t$ under which the rule fires. □

(a) is a real if minor defect: a top-up lands 3% short, and the trigger persists so the next tick tops up again. (b) needs its direction named. The rule errs *on the side of solvency*, and only in the region where it is used — but it buys that by over-reducing the short, which leaves the treasury under-hedged by $(k_e - k_g)N/P$ shares, 1.9% of the position at the worked point. By Theorem 4 residual delta is the headline correctness requirement, so "conservative" here means conservative against liquidation and adverse to neutrality: two different safeties pulling opposite ways. The exposure is transient — the next tick sees a drift outside the band and re-hedges — but it is the correct reading of the trade-off, not an unambiguous win. We note also that describing the Lighter model as $d \approx 1/\lambda - \mu$ for isolated leverage λ is optimistic by the factor $1/(1+\mu)$: the exact form is $(1/\lambda - \mu)/(1+\mu)$, a 0.6 percentage-point overstatement at $d = 0.2$.

6. Custody: what the automation can and cannot do

The keeper is an unattended process holding an EOA key and a trade-only venue key. The design's claim is not that it is trustworthy but that its authority is bounded. That is a reachability statement about immutable storage, and it is provable.

> **Theorem 12 (reachable set).** Let $\mathcal{A} = \{\mathcal{S}, \text{controller}, \text{Lighter account of } \mathcal{S}\}$. Every keeper-invocable path *settles* in $\mathcal{A}$, whose members are fixed at construction and all owned by $\mathcal{S}$. The one exception is deliberate and is the subject of Theorem 14: `allocate`'s swap leg pays the pool, whose counterparties are of course outside $\mathcal{A}$. So the honest statement is: *no keeper action can choose a destination*; the only value that leaves is what a swap at an oracle-floored price gives up.
>
> *Proof.* By enumeration of the router's six non-view externals (checked against the compiled ABI: there is no `receive`, no `fallback` and no payable function, so there is no ETH path at all). `allocate` pulls from $\mathcal{S}$ and pays out to (i) the swap router, transiently, under an oracle-derived `amountOutMinimum`, (ii) `zkLighter.deposit(to = safe, …)`, and (iii) the controller; `topUpMargin` pays only (ii); `fundLp` pays only (iii); `rescue` pays only $\mathcal{S}$ and is permissionless. `safe`, `controller`, `zkLighter` and `swapRouter` are `immutable`, and the constructor rejects a `swapRouter` the controller does not already trust and a controller whose `safe()` disagrees. The constructor also grants unlimited, permanent approvals to those two addresses (`forceApprove(swapRouter, max)` on the stock, `forceApprove(zkLighter, max)` on USDG) — worth naming in a reachability argument, since it is why pinning `swapRouter` to the controller's own is the load-bearing constructor check — the source calls it "the one deploy-time address that could take everything", and that caveat be-

longs here rather than in the contract alone. One of the four fixed venue values is *not* validated at construction: the asset index is checked against `zkLighter.tokenToAssetIndex(usdg)`, but `lighterRoute` is stored unchecked. It reads 0 on chain and no keeper action can change it, so the theorem holds; but "fixed at construction" is doing work here, and one of those fixings was never verified against the venue. The keeper cannot call `setKeeper`/`setPaused` (`onlySafe`). Lighter withdrawals initiated with the trade-only key are paid by the protocol to the account owner, $\mathcal{S}$, and do not transit the router. □

Lemma 13 (the fixed-window cap is exactly 2×). With a cap C enforced per fixed 24 h window, the notional moved in an arbitrary interval of length 24 h is at most $2C$, and this bound is attained.

Proof. An interval of length T intersects at most $\lceil T/T \rceil + 1 = 2$ fixed windows of length T, each contributing at most C. Attained by spending C immediately before a boundary and C immediately after. □

Theorem 14 (bounded leakage under full keeper compromise). Selling is priced against the oracle, not the pool: `allocate` enforces $\text{out} \geq (1-\theta) \cdot$ (oracle value of the sold stock). Hence value lost per allocation is at most θ of the *sold* notional. The cap is charged the whole allocation while only m of it reaches the swap, so per 24 h the leak is at most

$$2\,C\,m\,\theta \;=\; 2 \times 20{,}000 \times 0.40 \times 0.0105 \;=\; \$168,$$

with $2C\theta = \$420$ as the cruder bound that ignores the split. `topUpMargin` and `fundLp` consume cap but move USDG without a swap, so they leak nothing at all. A compromised keeper's total reach is that $168, plus the equity standing in the Lighter account (which it can lose by trading), and nothing else.

Proof. Theorem 12 bounds destinations; the `minOut` check bounds the price at which the only value-destroying step can execute; Lemma 13 bounds throughput. The Lighter term is not bounded by the router because the cap constrains inflow rate, not the standing balance — which is why the standing balance is sized to the hedge and monitored. □

The practical reading of Theorem 14 is that the deployed $20,000 cap converts keeper compromise from a solvency event into a $168/day nuisance plus the margin balance. The bound is linear in the cap, which is the only thing raising it from $2,000 on 2026-09-05 moved: *solvency is untouched*, because Corollary 9 makes the liquidation distance independent of how much flows through. Note what it does *not* protect: the router shares a Safe with an unrelated GME line, so `fundLp`/`topUpMargin` can move that line's USDG within the cap. This is why the first cap is small; it is a shared-custody bound, not a redundant one.

Remark 15 (residual signature set, as deployed). The claim "nobody signs in the loop" is true of the *allocate-and-hedge* loop and false of everything else, and the honest version is longer than it first appears. Every mutator on the controller is `onlySafe`: `recenter`, `recenterAsOwner`, `collect`, `close` and `sweepToSafe`. Add `setKeeper`/`setPaused` on the router, `FeeSink.migrateRecipient`, and a secure withdrawal of margin (an L1 signature the trade-only key cannot produce), and the deployed set is nine operations, not four.

Exactly one of them is designed to leave that set. Because a Zodiac Roles modifier executes with the Safe as `msg.sender`, the contract cannot tell the keeper from its owner, so the separation is by *selector*: a role scoped to `recenter(uint24)` alone opens that door to the keeper while `recenterAsOwner` — the same operation without the cooldown — stays the Safe's. **That grant has not been made.** The modifier is enabled on the Safe and is its only module, but the keeper holds no membership in any role (`NoMembership()` on chain for both controllers, checked 2026-09-05), so today every re-centre is a signature. Naming this is part of the claim rather than an exception to it: the automation this paper bounds is the allocation and the hedge, and the position's own lifecycle is still manual.

7. Accrual: what the token gets, and what it does not

Pons v2 has no holder distributor: the creator tax goes entirely to the fee recipient. The treasury therefore cannot and does not pay dividends, and the token carries no claim on treasury assets. The only channel from treasury performance to token holders is buy-and-burn. Two properties of that channel are provable, and they are the honest content of "回购销毁".

Theorem 16 (the buyback cannot deplete principal). Let K be treasury capital, ρ its realised yield rate, $\dot{\mathcal{T}}$ turnover rate, and let the buyback spend $B \le \max(0, \rho K)$. Then

$$\dot{K} \;=\; \tau\dot{\mathcal{T}} + \rho K - B \;\ge\; \tau\dot{\mathcal{T}} + \min(0, \rho K),$$

i.e. the buyback removes nothing the yield did not first produce. With $\dot{\mathcal{T}}$ constant, capital is bracketed by the two extremes of the policy,

$$\tau\dot{\mathcal{T}}\, t \;\le\; K(t) \;\le\; \frac{\tau\dot{\mathcal{T}}}{\rho}\left(e^{\rho t} - 1\right) \qquad (\rho > 0),$$

the lower bound at $B = \rho K$ (spend all of it: capital then grows linearly on tax alone) and the upper at $B = 0$ (compound everything).

Proof. Substituting $B = \rho K$ gives $\dot{K} = \tau\dot{\mathcal{T}}$, hence the linear lower bound; $B = 0$ gives $\dot{K} = \tau\dot{\mathcal{T}} + \rho K$, whose solution from $K(0) = 0$ is the exponential upper bound. Any admissible B lies between. □

What this does not say. An earlier draft claimed capital is non-decreasing on every path. It is not: when $\rho < 0$ the constraint forces $B = 0$, and $\dot{K} = \tau\dot{\mathcal{T}} + \rho K$ is negative once $|\rho K| > \tau\dot{\mathcal{T}}$ — which is precisely the regime Criterion 7 warns about, the gamma rent outrunning the fees. The theorem is about the *buyback*, not about the treasury: a losing market maker still loses money, it just does not do so by buying its own token back.

The contrast is with a fixed-rate buyback or dividend promise, under which B is independent of ρ and $\dot{K}$ can be negative — the standard mechanism by which "revenue-sharing" tokens consume the treasury that justified them. The constraint $B \le \rho K$ is what makes the difference, and it is a policy that must be enforced in the (as yet unwritten) executor, not a property of the contracts already deployed.

Theorem 17 (buy-and-burn strictly raises the capitulation floor). Let the graduated pool hold Q quote and T_p tokens with $k = QT_p$, and let T_c be tokens held outside the pool. Define the capitulation floor $P_\downarrow$ as the marginal price after every circulating token is sold into the pool: $P_\downarrow = k/(T_p + T_c)^2$. A treasury buy of B quote followed by burning the tokens received strictly increases $P_\downarrow$.

Proof. The buy leaves k invariant (and increases it by the retained fee), while replacing T_p by $T_p - \Delta$ with $\Delta > 0$; T_c is unchanged, since the acquired tokens are burnt rather than resold. Hence $P_\downarrow' = k'/(T_p - \Delta + T_c)^2 \ge k/(T_p - \Delta + T_c)^2 > k/(T_p + T_c)^2 = P_\downarrow$. □

This is a weaker and more defensible statement than "the burn supports the price". It says the *worst case* improves monotonically: the price that would obtain if every holder capitulated is strictly higher after each burn, and permanently so, because the liquidity is locked and the supply is gone.

7.1 Premium capture is the mirror image of the gamma rent

When the meme's own flow, or a closed underlying market, pushes the pool a fraction δ above fair value, the treasury can sell inventory into the pool and buy the same quantity back on the perpetual.

Proposition 18. The operation is delta-neutral by construction and its P&L is path-independent: selling ΔS shares into the pool at average $\bar{P}_{\text{pool}}$ and buying ΔS on the perp at P changes delta by $-\Delta S + \Delta S = 0$ and realises $\Delta S(\bar{P}_{\text{pool}} - P)$. Flattening a premium δ completely requires $x_q(1 - 1/\sqrt{1+\delta})/P$ shares of inventory and captures

$$\Pi_\delta \;=\; x_q\,\frac{(\sqrt{1+\delta} - 1)^2}{\sqrt{1+\delta}} \;\approx\; \tfrac{1}{4} x_q\, \delta^2.$$

Proof. Delta: the spot sale removes ΔS of exposure and the perp buy adds ΔS. P&L is the price difference times size, with no dependence on the path taken between the two legs (only on their simultaneity, which is an execution risk, not a price risk). For the capture, integrate the v3 curve from $P_1 = (1+\delta)P$ down to P: quote received is $L(\sqrt{P_1} - \sqrt{P})$ against shares $L(P^{-1/2} - P_1^{-1/2})$, and the excess over fair value is $L(\sqrt{P_1} - \sqrt{P})^2/\sqrt{P_1}$; substitute $x_q = L\sqrt{P}$. □

Note the form: $\frac{1}{4}x_q\delta^2$ against Theorem 5's $\frac{1}{4}x_q\sigma^2$. Passive liquidity pays the gamma rent; inventory deployed *as the informed side* collects the same functional quantity. Premium days are the days the treasury stops being a market maker and becomes the arbitrageur, and that is the economic reason for holding inventory at all rather than only margin. In the practical regime our inventory is far below the flattening requirement (Table 3), so the capture is linear, $\approx I \cdot \delta$, and the treasury takes a slice of the premium rather than closing it.

8. Reflexivity: the pairing-eligibility criterion

The design's most interesting failure mode is self-inflicted. Demand for the meme induces buying of the pairing asset, which pushes the pool away from the oracle; past a threshold that dislocation is negative for a hedged LP (measured at −38%/day on a pumped thin name, HIMS) and, worse, it trips the router's own deviation gate and stops the allocation exactly when fees are largest (Theorem 4).

> **Criterion 19 (pairing eligibility).** A net quote inflow F into a v3 pool moves the price by $(1+F/x_q)^2-1$. Requiring the induced move to stay inside the router's gate δ_{gate} gives
>
> $$F \;\le\; x_q\left(\sqrt{1+\delta_{\text{gate}}}-1\right) \;\approx\; \tfrac{1}{2}x_q\,\delta_{\text{gate}}.$$
>
> A pairing asset is eligible if the meme's plausible peak flow is well inside this bound.
>
> *Proof.* In v3, $\sqrt{P}$ is affine in quote reserves with slope $1/L$: $\Delta\sqrt{P} = F/L$, so $P'/P = (1+F/(L\sqrt{P}))^2$. Invert for F. □

At the measured NVDA/USDG 0.05% depth, $F \le \$536{,}532$: a memecoin would need half a million dollars of net flow through that pool, within one gate window, to gate its own treasury. Note the direction. Demand for the meme buys the pairing asset, which pushes the pool *up*; under the deployed symmetric gate that is exactly what trips it, which is the reflexivity this criterion prices. Under the proposed one-sided gate (§4) an upward dislocation no longer blocks the sale at all, and this criterion stops binding in the direction a successful launch pushes — it would then bind only on a collapse. That is a real improvement and it is worth being explicit that the criterion below is stated against the gate as deployed today. The same computation on a \$9.5k pool gives $F \le \$785$, which any launch clears in its first minute. This is the quantitative content of the decision to pair against NVDA and to reject a thin-pool ticker for the second launch, and it is a criterion that can be evaluated on any candidate before committing: the pairing asset's depth must exceed roughly $2F_{\text{peak}}/\delta_{\text{gate}}$.

9. Measurements

All figures below are computed by `docs/paper/verify_numbers.py` from a 92.2-hour measurement window ending 2026-09-05. That window is *cumulative, not rolling* — `lp_calc.py` reads every row collected so far — so rerunning the command later lengthens the window as well as refreshing it, and a reader who gets different numbers cannot tell the two apart without checking the hours. The figures are a snapshot as of that timestamp (`data/pool_flow.csv`, `data/basis.csv`, `data/swaps_*.csv` via `lp_calc.py`). The example position is \$20,000 in a $\pm 3\%$ band, the scale the treasury will actually operate at in its first months.

quantity	value	source
spot, daily σ of log returns	`$231.18, 1.10%`	measured
pool volume $\dot{V}$	`$38.44M/day`	measured
pool depth x_q^{pool}	`$214.88M`	from L_p, Thm. 5
pool fee revenue $\varphi\dot{V}$	`$19,220/day`	Criterion 7
pool LVR $\frac{1}{4}\sigma^2 x_q^{\text{pool}}$	`$6,500/day`	Thm. 5
entry test $\varphi\dot{V}/\ell_{\text{pool}}$	**`2.96 > 1 ✔`**	Criterion 7
profit-maximising share $s^\star$	`41.9% (≈$4.7M)`	Criterion 7

our $20k position at ±3%:		
depth x_q, fee share s	$661,629, 0.307%	§4.2
capital deployed in the LP (quote leg $10,000)	$19,707	App. A.2
fees	+$59.00/day	measured flow
funding received on the short	+$1.50/day	measured median
gamma rent, σ-based (Thm. 5)	−$20.01/day	model
markout, measured against the 5-min mark	+$34.22/day	measured, $148M sample
hedge execution + re-centring + gas	−$1.90/day	measured
capital committed (LP + margin)	$32,844 ($13,138 margin)	$C/(1-m)$
net, σ-based (conservative)	**+$38.58/day = 42.9% APR**	on committed capital
… on the LP leg alone	71.5% APR	—
net, markout-based	+$92.82/day = 103.1% APR	on committed capital
expected in-range spell before a re-centre	7.4 days (49/yr)	§9
breakeven pool volume, σ-based	$13.30M/day	—
revenue and safety:		
creator tax to graduation	1.448 NVDA ≈ $335	Cor. 2
creator tax per $1M of turnover	$33,300	Thm. 1
liquidation distance, $\rho = \frac{1}{2}$ / 1	+126.5% / +61.8%	Thm. 8
minimum viable m for a 15.8% gap	16.16% (deployed 40%)	Cor. 10
max daily leakage, compromised keeper	$168 ($420 over-bound)	Thm. 14
inventory to flatten a 1% premium	4,613 NVDA ($1.07M)	Prop. 18
flow that gates the treasury	$536,532	Criterion 19

Table 3. Measured and derived parameters. The two LVR estimates disagree in sign, which is the most important line in the table; see below.

Two things moved at once. The σ-based net in this table is $38.58/day where an earlier draft had $24.35, and that draft also lacked the cost line. The two effects pull opposite ways and are worth separating: re-measuring on the longer window took it from $24.35 to $40.49 (in-range liquidity fell while volume held, so our share rose), and adding the execution, re-centring and gas line then took $1.90 back off. The cost line is not free; it is simply smaller than the re-measurement.

On the denominator. The LP leg is $(1-m)$ of an allocated cohort, so holding $20,000 in the pool commits $32,844, the remaining $13,138 sitting as Lighter margin. Reporting the return on the LP leg alone would flatter it by $1/(1-m)$ — and would be internally inconsistent, since the funding in the numerator is income earned *on that margin*. Table 3 therefore leads with the return on committed capital and gives the LP-only figure beneath it. Both are in-range rates: the fee share and the gamma rent alike stop applying when price leaves the band, which at $\sigma = 1.10\%$ and $\pm 3\%$ happens about every 7.4 days, and each exit is a re-centre that realises divergence and pays two swaps. The measured re-centring cost is in the table; the divergence it realises is already inside the markout and the gamma rent, not a fourth term.

On the sign disagreement. The σ-based gamma rent says the position bleeds $20/day to price dynamics; the realised markout against the perp mark says it *gains* $34.22/day, over a $148M / 164k-swap sample. Both cannot be right about the same quantity. Under a martingale mark they estimate the same thing (§4.1), so the disagreement localises the failure of that assumption: either the Robinhood Chain pool leads the Hyperliquid `xyz` mark at the five-minute horizon (making the markout flattering), or the five-minute horizon truncates informed flow that would show up at longer horizons, or both. We therefore report the σ-based figure as the number to plan against — 43% APR on com-

mitted capital, not 103% — and treat the markout as evidence that the flow in this pool is unusually retail. The design is profitable under either estimate; that robustness, rather than either point estimate, is the claim.

10. What is not proved

The theorems above are all conditional on the venue and the contracts behaving as specified, and none of them addresses the load-bearing economic input. In descending order of how much they should worry a reader:

1. **Turnover is exogenous.** Theorem 1 makes revenue a clean function of $\mathcal{T}$, and nothing in this design produces $\mathcal{T}$. If the token does not trade, every property proved here holds exactly and the treasury stays at $334. This is the whole risk of the venture, and it is not a mathematical one.
2. **The yield term is a measurement, not a theorem.** Criterion 7 passes at 3.0× today on one pool over 92 hours — and it read 2.0× on the same pool thirteen hours earlier, because in-range liquidity fell 28% while volume held. A criterion that moves by half in half a day is not a fact about the design, it is a measurement with a shelf life. Fee tiers get undercut, volume migrates, and σ is not stationary. Re-measure before sizing up; the entry test is cheap enough to run every tick, which is what the keeper now does.
3. **Basis risk survives a correct hedge.** Theorem 3 zeroes direction, not the spread between the three prices the legs are marked at. A pool that dislocates from the perpetual mark moves the LP leg against the short with no delta error to detect; the weekend is the standing example, when the oracle is frozen, the pool trades on and the mark follows a different reference. It is bounded by the size of the basis rather than by the size of the position, which is why it is small in normal weeks and is not small on a premium weekend — the same episode Proposition 18 is designed to earn from.
4. **The custody bound covers the router, not the future keeper.** Theorem 12 enumerates the *router's* external functions. The roadmap delegates `recenter(uint24)` to the keeper through a Roles modifier, and a re-centre performs up to two pool swaps sized off the whole position value, bounded by the controller's own tolerance and its `minRecenterInterval` but *not* by `dailyCapUsdg`. Once that grant is signed, worst-case extraction stops scaling with C and starts scaling with position size divided by re-centre cadence. The bound in §6 is a statement about the authority set deployed today (Rem. 15), and it will need restating the day that changes.
5. **A global pause of the equity token is the one unhedgeable event.** Tokenised equities on this chain can be paused by the issuer. Under a pause the LP and Safe legs are frozen while the perpetual keeps trading: value neutrality (Thm. 3) still holds on paper, but the hedge can no longer be adjusted or unwound, so what was a delta-neutral book becomes a naked short whose loss is bounded only by Theorem 8's distance. Nothing in the design mitigates this; only position sizing does.
6. **Oracle behaviour when the underlying market is closed.** The deviation gate compares the pool to a Chainlink price that is fresh but references a market that is not trading. A weekend premium is therefore indistinguishable, to the gate, from a manipulation — correctly, it refuses to sell — but the consequence is that stock accumulates unallocated in the Safe over exactly those windows (Thm. 4), and only the keeper's exposure sum keeps it hedged.
7. **Revenue arrives with an uncontrolled lag.** Creator fees reach the claimable escrow only after the launchpad operator sweeps them. Theorem 1 gives the total; it says nothing about when, and the treasury's reporting must show the unswept portion separately.
8. **Pre-positioning inside the gate.** Theorem 14 bounds extraction at θ of the sold leg: an adversary can push the pool to the inside edge of the gate and wait for a scheduled allocation. $168/day at the current cap, growing linearly with it. Randomising allocation timing reduces it; nothing eliminates it short of pricing the swap on the oracle alone.
9. **Shared custody.** The router's cap is shared with an unrelated line's USDG in the same Safe (§6). The bound is correct but the blast radius is larger than one product.
0. **Jurisdiction and terms.** The perpetual venue geo-fences order placement; the keeper's execution locus is a compliance fact, not an implementation detail. The token carries no claim on any treasury asset, and nothing here should be read as offering one.

11. Verdict

The design is mathematically sound in a specific and limited sense, and it is worth being precise about which:

Sound. Delta neutrality is an exact identity (Thm. 3), which is unusual — most hedged-LP constructions are neutral only to a linearisation — and the identity also tells the keeper what to read. The solvency bound is closed-form and scale-free (Thm. 8–9), so the safety argument does not need re-litigating as the treasury grows; the deployed 40% carries 2.5× the minimum the venue's own worst gap it is sized against requires. The custody argument is a reachability proof over immutable storage with a $420/day leakage bound (Thm. 12–14), which is a genuinely different kind of assurance from "the keeper is careful". The accrual channel is monotone and cannot eat principal (Thm. 16–17). The yield condition passes on measurement by 3× today (2× half a day earlier), and its optimal-share formula gives a growth path rather than a guess.

Not sound as often stated. "Delta-neutral" is routinely heard as "risk-free": it is not, and Theorem 5 is the precise reason — neutrality removes variance and leaves a deterministic gamma rent that fees must outrun. "The LP earns the treasury's yield" understates by two orders of magnitude in the wrong direction: the position earns $39–93/day on $33k committed while a single $1M day of meme turnover pays $33,300, so the hedged LP is a mechanism for making inventory non-directional and mildly yield-bearing, not the engine. And "nobody signs" is true only of the steady state (Rem. 15).

The honest one-line summary. What is proved is that this treasury cannot become a levered directional bet on NVDA, cannot be liquidated by anything short of a +62% overnight gap, cannot be drained by its own automation beyond $168/day, and cannot have its principal spent by its own buyback. What is not proved — and cannot be — is that anyone will trade the token. The mathematics makes the treasury safe to grow; it does not make it grow.

Appendix A. Derivations used above

A.1 v3 holdings and value. For liquidity L on $[P_a, P_b]$ with $P \in [P_a, P_b]$, the standard amounts are $x(P) = L(P^{-1/2} - P_b^{-1/2})$ shares and $y(P) = L(\sqrt{P} - \sqrt{P_a})$ quote, whence $V(P) = y + xP = L(2\sqrt{P} - \sqrt{P_a} - P/\sqrt{P_b})$, $V'(P) = x(P)$, $V''(P) = -L/(2P^{3/2})$. The controller's `_amountsForLiquidity` implements the same three-branch formula in fixed point.

A.2 Capital and depth for a symmetric band. The two legs of a price-symmetric band are not equal: $y = x_q(1 - \sqrt{1-r})$ against $xP = x_q(1 - 1/\sqrt{1+r})$. So "the quote leg is $C/2$" and "the position deploys C" cannot both be true, and it is worth saying which one the numbers use. Total capital is

$$C \;=\; x_q\left(2 - \sqrt{1-r} - \tfrac{1}{\sqrt{1+r}}\right), \qquad \text{so} \qquad x_q \;=\; \frac{C}{2 - \sqrt{1-r} - 1/\sqrt{1+r}},$$

a concentration multiplier of 33.6× at $r = 3\%$. The figures in §9 come from `lp_calc.py`, whose $\min(L_x, L_y)$ binds on the quote leg, so its L is the one with a $10,000 quote leg: $x_q = \$661,629$, and the position it describes deploys $\$19,707$, not $20,000. We keep that L and report the capital it actually represents rather than rescaling to a round number — the alternative is a table whose numbers no longer match the command that produced them. Theorem 6 is the statement that this multiplier does not appear in the sign of the edge.

A.3 Price impact. $\sqrt{P}$ is affine in the quote reserve with slope $1/L$, so a quote inflow F gives $P'/P = (1 + F/x_q)^2$; inverting yields Criterion 19, and the same relation with the sign reversed gives the shares required in Proposition 18.

A.4 Liquidation. Equity $M + q(P_0 - P)$ against maintenance $\mu q P$ gives $P_{\text{liq}} = (P_0 + M/q)/(1+\mu)$; with $M = mSP_0$ and $q = (1-m)\rho S$ the distance is $\big(m/((1-m)\rho) - \mu\big)/(1+\mu)$, free of S and P_0.