Designing a Decentralized Protocol for Hedging Concentrated Liquidity Positions

DLT2026 Workshop

S. Avanzo, L. Pennella, A. Norta, N. Yadav, A. Kumar, M. Singh, A. Sunday, G. Peraldo, C. Schifanella

DLT2026

University of Torino | Norta DeSyCo | Tallinn University | University of Pretoria | University of Luxembourg

Introduction: DEXs, AMMs & Concentrated Liquidity

- **Decentralized Exchanges (DEXs).** Trade tokens on-chain, peer-to-peer, with no intermediary.
- **Automated Market Makers (AMMs).** Most DEXs: an order book is replaced by a *liquidity pool* priced by a formula. **LPs** supply token pairs and earn fees; **traders** swap against the pool and pay them.
- **Uniform liquidity** (e.g. Uniswap v2) — constant product [2]:

$$x\,y = L^2$$

- **Concentrated liquidity (CLMMs)** (Uniswap v3, Orca) — capital confined to a range $[p_a, p_b]$, deepening liquidity near the price [2]:

$$\left(x + \frac{L}{\sqrt{p_b}}\right)\left(y + L\sqrt{p_a}\right) = L^2$$

$\Rightarrow$ higher fees, but stronger *impermanent loss*.

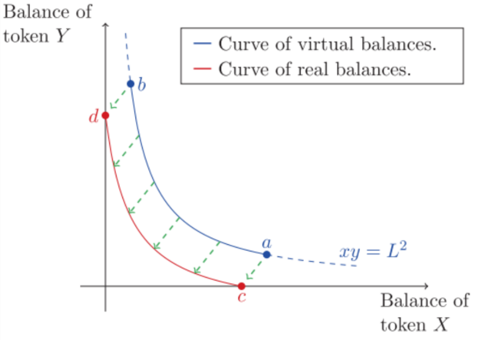

Virtual ($x\,y = L^2$) vs. real reserves of a CL position [2].

Motivation & Research Questions

Problem. Liquidity Providers in Concentrated-Liquidity Market Makers (CLMMs) suffer *impermanent loss*, especially for tight ranges $\Rightarrow$ stronger depreciation risk of LP positions [1].

Gap in the state of the art.

- Existing hedges (options, perpetuals, dynamic rebalancing) depend on *liquid derivative markets* that are absent for most token pairs [4, 5, 6].
- No *protocol-level* mechanism for direct, range-specific risk transfer.

Main RQ. *How to design a decentralized protocol to enable hedging of the equity of liquidity positions in CLMMs?*

- **RQ1.** *What is the set of requirements that enables a decentralized protocol to support the transfer of liquidity-position depreciation risk between participants?*
- **RQ2.** *What is a suitable model for pricing the risk fairly, ensuring sustainable participation of the involved parties?*

Methodology & Contributions

Methodology. Design Science Research [17]; protocol requirements elicited with the *Trusted DApp Modeling* (T-DM) method [18, 19].

Contributions

- **Protocol conceptual design (RQ1).** A T-DM goal model deriving the on-chain, functional, and quality requirements for protocol-level transfer of concentrated-liquidity risk.
- **Pricing model (RQ2).** A premium pricing model for a capped bilateral hedge; a theorem proving the protocol's *Value Neutrality*; a breakeven heuristic deciding, from live on-chain data, when hedging pays off per pool.
- **Implementation.** A proof of concept, one module per goal-model sub-goal, demonstrating the system's financial feasibility and utility; the protocol's own contracts are mocked off-chain, while it opens and reads real Orca Whirlpool positions on Solana mainnet.
- **Evaluation.** A 52-week SOL/USDC backtest (ranges $\pm 5/7.5/10\%$): reductions in volatility, max drawdown, and tail risk; Value Neutrality verified. A live mainnet test confirms end-to-end execution on a real pool.

Background: Notation and Core Definitions

- A Liquidity Provider (**LP**) earns trading fees on a concentrated-liquidity pool, but the position loses value when the price moves.
- Our protocol lets the LP transfer that risk to a risk-taker (**RT**) pool in exchange for a weekly premium.
- Recurring symbols are summarised below; contract- and pricing-specific notation appears where first used.

Actors & timing

- **LP** — liquidity provider, *buys* protection
- **RT** — risk-taker pool, *sells* protection
- w — one weekly settlement cycle

Position model

- S_t, S_0 — token price now / at issuance
- σ — price volatility
- $[p_\ell, p_u]$ — chosen liquidity price range
- L — position liquidity (its size)
- $V(S)$ — value of the LP position when the price is S

Background: Position Value & Impermanent Loss

CL position value $V(S)$ [2] is *piecewise*:

- **Inside** $[p_\ell, p_u]$: holds both tokens; inventory rebalances against the move (sheds the winner, accumulates the loser) $\Rightarrow$ V smooth and *strictly concave*.
- **Outside**: fully converted — linear below p_ℓ, flat above p_u; the position stops tracking the price.

Impermanent loss (IL)

- $\mathrm{IL}(S) = H(S) - V(S) \geq 0$: the gap versus simply *holding* the initial tokens (H).
- Zero at S_0; grows the further the price drifts (see chart).
- The implicit price of earning fees — the risk this protocol lets the LP transfer away [3].

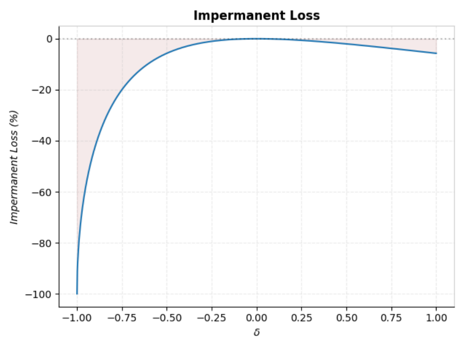

Figure 5: Impermanent Loss in Uniform Liquidity AMM as a function of change in price (δ) [4].

Requirements (RQ1): T-DM Goal Model

Root goal: *enable risk transfer for LPs.*

Global Quality Goals

- *Economically sustainable*: keep breakeven wedge $r^{\star}(\mathcal{S}) - r_u$ small.
- *Incentive-aligning*: RT compensated by premium *and* fee-share.

Six functional sub-goals:

1. *Retrieve and dispatch market data*
2. *Manage certificate lifecycle*
3. *Maintain protocol operation*
4. *Quote premiums*
5. *Manage escrowed Whirlpool position lifecycle*
6. *Maintain protection pool and solvency controls*

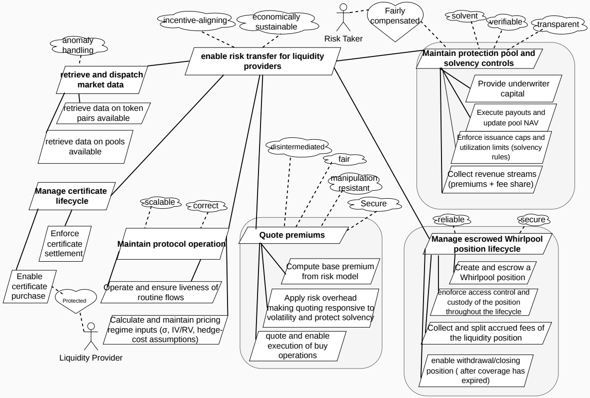

Goal model of the Liquidity Hedge protocol.

The Liquidity Hedge Contract: Three Legs

An LP buys a tokenized *certificate* from a Risk Taker (RT) pool. Each weekly cycle has three cash-flow *legs*:

- **Premium leg** P_w — paid upfront by LP, split RT / treasury φ.
- **Price-contingent leg** Π_w — settlement, capped by range boundaries.
- **Fee-split leg** $y\,F_w$ — share of LP's realised trading fees to RT.

Range-clamped payoff (LP perspective):

$$\Pi(S_T) = V(S_0) - V\big(\mathrm{clamp}(S_T, p_\ell, p_u)\big) = \begin{cases} +\mathrm{Cap}_{\downarrow} & S_T \le p_\ell, \\ V(S_0) - V(S_T) & p_\ell < S_T < p_u, \\ -\mathrm{Cap}_{\uparrow} & S_T \ge p_u. \end{cases}$$

S_T: settlement price • $\mathrm{clamp}(\cdot)$ pins S_T to the nearest range edge • $\mathrm{Cap}_{\downarrow}, \mathrm{Cap}_{\uparrow}$: fixed max down/up payoff outside $[p_\ell, p_u]$.

Convexity wedge. For symmetric ranges, concavity of V gives $\mathrm{Cap}_{\downarrow} > \mathrm{Cap}_{\uparrow} \Rightarrow$ *strictly positive* fair value. The same boundaries cap the RT pool's *maximum loss* ($\mathrm{Cap}_{\downarrow}$), bounding downside risk and supporting solvency.

Pricing the Risk (RQ2): Canonical Premium

Fair value. FV_w is what the hedge would cost in a frictionless market — its risk-neutral expected payoff under GBM [11, 12]:

$$\mathrm{FV}_w := \mathbb{E}_{\mathbb{Q}}\big[\,\Pi(S_T)\,\big] \;>\; 0 \quad (\sigma > 0,\ S_0 \in (p_\ell, p_u)).$$

Canonical premium. What the LP actually pays each week: the fair value loaded for risk, minus the compensation the RT already earns from the fee split.

$$P_w \;=\; \max\Big(P_{\mathrm{floor}},\ \mathrm{FV}_w \cdot m_{\mathrm{vol},w} \;-\; y\,F_w(r)\Big).$$

- **Implied risk compensation parameter** $m_{\mathrm{vol},w} = \max(\bar{m},\ \mathrm{IV}_w/\mathrm{RV}_w)$: the implied-to-realised volatility ratio charges more when the market expects more turbulence than recently seen. The governance parameter $\bar{m}$ floors this multiplier, ensuring a minimum markup.
- **Fee discount** $-\,y\,F_w(r)$: the LP already hands a share y of its trading fees F_w to the RT, so the premium is cut up front by this expected amount — avoiding double payment for the same compensation.
- **Floor** P_{floor}: a minimum premium that keeps underwriting worthwhile when volatility is very low.

Governance parameters: P_{floor}, y, $\bar{m}$, plus the treasury fee $\varphi \in [0, 1)$ levied on every premium.

Value Neutrality & the Breakeven Wedge

Value Neutrality. Adding the LP's and the RT's weekly P&L, the settlement and fee legs cancel:

$$\mathrm{LP}_w(r) + \mathrm{RT}_w(r) = U_w(r) - \varphi\, P_w.$$

So the hedge only *redistributes* U_w; the *sole* leakage is the treasury's cut φP_w of each premium.

Breakeven wedge. What extra fee yield must the position earn to absorb that leak? Weekly fees are $F_w = 7\, V_w\, r$ (value $\times$ 7 days $\times$ daily yield), so each extra unit of daily yield adds $7 \sum_w V_w$ over the horizon. Dividing the dollar leak $\varphi \sum_w P_w$ by this fee base gives it:

$$\boxed{r^\star - r_u \;=\; \frac{\varphi \sum_w P_w}{7 \sum_w V_w}}$$

i.e. $r^\star$ (smallest yield at which *both* LP and RT break even) sits exactly this far above the unhedged r_u.

Implications

- It is the *entire* cost of hedging, fixed by the single governance parameter φ; as $\varphi \to 0$ the wedge vanishes and the hedge is free in aggregate.
- *Viability becomes easily decidable:* hedging pays off in any pool whose live fee yield exceeds $r^\star = r_u + \text{wedge}$.

Evaluation: Risk Reduction (52 wk SOL/USDC)

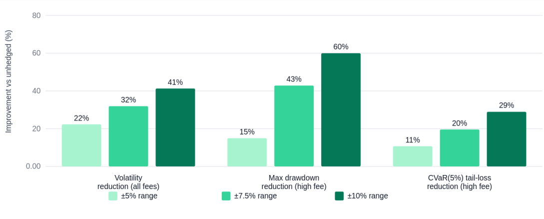

Volatility reduction, **MaxDD reduction**, $\mathrm{CVaR}(5\%)$ **reduction (high-fee tier)**.

[13, 15]

Setup. Backtest on 52 wk SOL/USDC (Birdeye); symmetric ranges $\pm 5\%, \pm 7.5\%, \pm 10\%$; σ_{ann}=57.3%; y=10%, φ=1.5%, $\bar{m}$=1.05. LP fee tiers 0.10/0.25/0.45% per day.

Volatility reduction. 22.4% $\rightarrow$ 41.3% as range widens.

Maximum drawdown reduction (by range width):

high-fee tier: 14.9%, 42.8%, 60.0%

low-fee tier: 2.2%, 10.3%, 15.0%

Benefit scales with fee income: premiums partly offset it when fees are low.

$\mathrm{CVaR}(5\%)$ **improvement:** $10.7\%, 19.6\%, 28.9\%$.

Value Neutrality verified empirically: residual $|\mathrm{LP}_w + \mathrm{RT}_w - (U_w - \varphi P_w)| < 10^{-6}$ USDC. The residual is due to floating-point rounding, hence confirming the theorem.

Mean-Based Metrics & Breakeven Viability

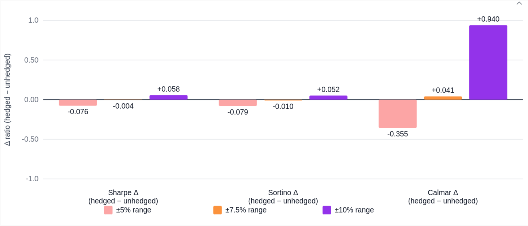

Sharpe / Sortino / Calmar Δ = hedged − unhedged. [16, 14]

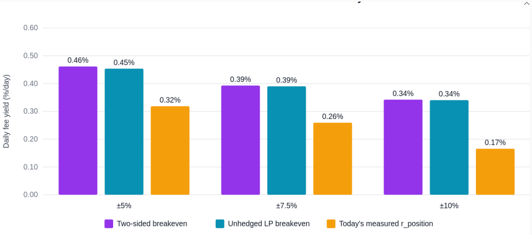

Two-sided breakeven $r^\star$, unhedged r_u, and live r_{position}.

- Sharpe/Sortino/Calmar Δ turn positive at $\pm 10\%$: $\{+0.058, +0.05, +0.94\}$ $\Rightarrow$ wide ranges make hedging benefits dominate the premium cost.
- Empirical wedges $r^\star - r_u \in \{0.8, 0.3, 0.2\}$ bps/day match the closed-form prediction.
- Live SOL/USDC pool yield r_{position} falls below $r^\star$ at every width: viability is *pool-specific*, decidable from two live-measurable quantities.

Conclusion

Contributions.

- **RQ1:** T-DM goal model deriving on-chain / functional / quality requirements for protocol-level CL risk transfer.
- **RQ2:** Closed-form pricing of a capped bilateral hedge; *Value Neutrality* theorem $\Rightarrow$ wedge driven by a single governance parameter (φ).
- **Deployment heuristic:** a live-data rule for *when* to hedge — deploy on any pool whose measured fee yield r_{position} clears the two-sided breakeven $r^\star$, so viability is decidable per pool from on-chain observables.
- **PoC & empirical results:** 52-week SOL/USDC backtest + live Orca Whirlpool settlement `github.com/SoweluAvanzo/LiquidityHedge`.

Limitations & future work.

- Empirical scope: extend beyond backtest + live tests with Monte-Carlo simulation.
- Benchmark vs. established hedging strategies (options, perpetuals).
- Move PoC from mocked on-chain components to a fully-fledged DApp.

References

- Hasbrouck, Rivera, Saleh. *Economic Model of DEX with Concentrated Liquidity.* Mgmt. Sci., 2025.
- Ottina, Steffensen, Kristensen. *Automated Market Makers.* Springer, 2023.
- Hashemseresht, Pourpouneh. *CL Analysis in Uniswap V3.* ACM CCS DeFi, 2022.
- Khakhar, Chen. *Delta Hedging LP Positions on AMMs.* arXiv:2208.03318, 2022.
- Schaller. *Hedging the Risks of LPs.* 2022.
- Zhang, Chen, Yang. *Adaptive LP in Uniswap V3 with DRL.* arXiv:2309.10129, 2023.
- Risk, Tung, Wang. *Pricing/Hedging in CFMMs.* arXiv, 2026.
- Kroo et al. *Optimal Portfolios of LP Positions.* BCCA, 2024.
- Pennella et al. *Unified Framework of DeFi Derivatives.* arXiv:2512.19113, 2025.
- Del Monte et al. *Current Understanding of IL Risk.* BCRA, 2025.

References (II)

- Black, Scholes. *Pricing of Options.* JPE, 1973.
- Hull, Basu. *Options, Futures, and Other Derivatives.* Pearson, 2016.
- Magdon-Ismail, Atiya. *Maximum Drawdown.* Risk, 2004.
- Sortino, Van Der Meer. *Downside Risk.* JPM, 1991.
- Acerbi, Tasche. *Coherence of Expected Shortfall.* JBF, 2002.
- Sen. *Sharpe, Sortino, Calmar Comparative Study.* 2022.
- Hevner et al. *Design Science in IS Research.* MISQ, 2004.
- Udokwu. *T-DM Method.* PhD Thesis, LUT, 2022.
- Sterling, Taveter. *Art of Agent-Oriented Modeling.* MIT, 2009.
- Mohan. *AMMs & DEXes: A DeFi Primer.* Fin. Innovation, 2022.