Schur Damping for Perpetual Demand Lending Pools

Peter Cotton

June 2026

Abstract

Perpetual demand lending pools (PDLPs) must decide whether to operate as one pool or several. Chitra et al. (2025) resolve this with the *undamped* Schur complement of the asset covariance and a binary spectral threshold: merge the pools if a single condition number test passes, keep them apart otherwise. We observe that the merged and isolated pools are the two endpoints, $\gamma = 1$ and $\gamma = 0$, of the Schur-complementary damping family of Cotton (2024), and we propose occupying its interior. A *partially merged* pool hedges against $A - \gamma\, BC^{-1}B^{\mathsf{T}}$, with the coupling fraction γ set not by a threshold but by the statistical reliability of the cross-pool covariance block B. We give the closed-form $\gamma^{\star}$ implied by viewing the choice as a James–Stein shrinkage problem (Cotton, 2026). Because PDLPs live in an acute undersampled regime—short on-chain return histories, fast rebalancing, asset counts that rival or exceed the number of return samples—$\gamma^{\star}$ is pulled toward 0: the theory predicts pools should stay *mostly* split, merging only as cross-pool coupling proves both strong and well-sampled. The result is a graded, data-driven target-weight mechanism that degrades gracefully where the binary test flips discretely, and it supplies precisely the pool covariance that GMX V2's dynamic-pricing upgrade is noted to omit.

1 Introduction

Perpetual futures are the most liquid crypto markets, and a growing share trades on decentralized venues—GMX, Jupiter, Hyperliquid—that fund market-maker leverage through on-chain liquidity pools. Chitra et al. (2025) formalize this mechanism as a *perpetual demand lending pool* (PDLP): liquidity providers deposit assets into a pool with a target composition, traders borrow against it to open leveraged positions, and arbitrageurs keep the pool near target. The authors give conditions under which a PDLP's liquidity providers can be delta hedged, partially explaining why these pools have grown to billions of dollars in assets.

A recurring architectural question, raised explicitly in their Appendix B, is when several pools should be combined into one. GMX itself migrated from a single pool to a multi-pool system; the inverse question—when to re-aggregate—is decided in Chitra et al. (2025) by partitioning the asset covariance into blocks and comparing a pool's own covariance against its Schur complement. The comparison is binary: a spectral inequality either holds, and one merged pool delta hedges better, or it fails, and the pools should stay apart.

This note makes a single observation and develops its consequence. The two options being compared—the isolated pool and the merged pool—are not two unrelated architectures. They are the $\gamma = 0$ and $\gamma = 1$ endpoints of the damped Schur complement that interpolates hierarchical risk parity and minimum variance (Cotton, 2024, 2022). The interesting regime for a PDLP is the interior, $0 < \gamma < 1$: a pool that is *partially* merged, sharing risk across pools in proportion to how reliably that shared risk can be estimated. In the data-poor conditions typical of DeFi, that proportion is small, and the interior solution differs materially from either endpoint.

2 The pool-merging decision is the γ-bridge endpoints

We recall the setup of Chitra et al. (2025, App. B). A PDLP holds risky assets with covariance $\Sigma \succ 0$. To ask whether two pools should be one, partition the assets into a pool-1 block of size m_1 and a pool-2 block of size m_2 with disjoint supports, $m_1 + m_2 = n$, and write

$$\Sigma = \begin{bmatrix} A & B \\ B^\mathsf{T} & C \end{bmatrix}, \qquad A \in \mathbb{R}^{m_1 \times m_1},\ C \in \mathbb{R}^{m_2 \times m_2},\ B \in \mathbb{R}^{m_1 \times m_2}, \tag{1}$$

with A, C nonsingular and positive definite. Here A is pool 1's own covariance, C is pool 2's, and B is the *cross-pool* block. For mean-variance delta hedging the relevant objects are the conditional covariances, the Schur complements

$$\Sigma/C = A - B\,C^{-1}B^\mathsf{T}, \qquad \Sigma/A = C - B^\mathsf{T}A^{-1}B. \tag{2}$$

A merged pool incorporates information from both blocks and hedges against Σ/C (resp. Σ/A); an isolated pool hedges against A (resp. C) alone, as though the other pool did not exist. Chitra et al. (2025) prove a sufficient condition for one pool to be better than two.

Claim 1 (Chitra et al., 2025, Claim B.1)**.** *If* $\sigma_{\max}(A) < \sigma_{\min}(\Sigma/A)$ *and* $\sigma_{\max}(B^\mathsf{T}) < \sigma_{\min}(\Sigma/C)$ *then the single merged pool delivers better delta-hedged returns than the two isolated pools.*

The authors note these conditions are "similar to hierarchical risk-parity methods" (Cotton, 2024; López de Prado, 2016). The similarity is exact, and it is the point of departure for what follows: the isolated pool keeps only its diagonal block A, while the merged pool conditions fully through $A - BC^{-1}B^\mathsf{T}$. These are the two ends of one dial.

3 Damped pool merging

Following Cotton (2024), interpolate the coupling. For $\gamma \in [0, 1]$ define the *damped sub-covariance*

$$A(\gamma) \;=\; A - \gamma\,B\,C^{-1}B^\mathsf{T}, \qquad C(\gamma) \;=\; C - \gamma\,B^\mathsf{T}A^{-1}B, \tag{3}$$

and let a *partially merged* PDLP delta hedge pool 1 against $A(\gamma)$ and pool 2 against $C(\gamma)$.

Proposition 1 (Endpoints)**.** $A(0) = A$ *recovers the isolated pool, and* $A(1) = \Sigma/C$ *recovers the fully merged pool of Chitra et al. (2025). The single-pool-versus-two-pools decision of Claim 1 is the comparison of the two endpoints* $\gamma \in \{0, 1\}$.

The proposition is immediate from (3), but it reframes the problem. Rather than testing which endpoint is better, choose γ. And γ has a meaning: it is how much of the estimated cross-pool coupling to trust.

Choosing γ. The right γ is the reliability of the estimated coupling, a James–Stein / Wiener ratio of signal to signal-plus-noise rather than a style preference (Cotton, 2026). In the two-block case with conditional correlation ρ estimated from n return observations,

$$\gamma^\star \;=\; \frac{(n-2)\,\rho^2}{(n-2)\,\rho^2 + (1-\rho^2)}. \tag{4}$$

Strong, well-sampled cross-pool coupling sends $\gamma^\star \to 1$ and the merge deserves trust; weak or undersampled coupling sends $\gamma^\star \to 0$ and the divide-and-conquer split was right. The estimate $\gamma^\star$ tracks the validation-tuned optimum with no tuning, on portfolios, crypto correlation, and simulated processes (Cotton, 2024, 2026).

4 The undersampled regime is where DeFi lives

The binary test of Claim 1 silently assumes B is known. In a PDLP it is the least-known part of the model. The pools are disjoint by construction, so B is a cross-*block* covariance—the entries with the fewest joint observations and the most estimation noise. On-chain return histories are short, pools rebalance on the order of hours, and the number of pooled assets can rival or exceed the number of return samples n. Trusting B completely (merge) overfits noise into the hedge; discarding it (split) throws away genuine diversification. The reliability-weighted answer is to trust B in proportion to how well it is measured, which is exactly (4).

Small n pulls $\gamma^\star$ toward 0. The theory therefore predicts that PDLP pools should remain *mostly* split, merging incrementally only as the cross-pool coupling proves both strong (large ρ) and stable (large n). This is a graded, falsifiable prediction. The threshold test of Claim 1 cannot make it: it flips from "split" to "merge" the instant a spectral inequality is met and is silent on the long approach to that boundary, which is where a young, data-poor pool actually operates.

This is the same identity that drives Cotton (2026): the damping a portfolio needs when assets outnumber returns is, term for term, the damping a spatial pseudo-likelihood needs when stations outnumber observations. A perpetual demand lending pool is an unusually sharp instance of the outnumbered regime, which is why its architecture question lives in the interior of the dial rather than at its ends.

5 A reliability-weighted target-weight mechanism

Chitra et al. (2025) maintain a pool near a target portfolio through a target-weight mechanism, and close Appendix B by noting that GMX V2's dynamic-pricing upgrade "does not take into account pool asset covariance." Equation (3) supplies exactly that covariance, at a reliability-weighted amount. Concretely, a damped PDLP hedges pool 1 against $A(\gamma^\star)$ rather than against either A or Σ/C, and a target-weight mechanism can set inter-pool collateral sharing proportional to $\gamma^\star$:

$$\text{shared coupling} \;=\; \gamma^\star \cdot B\,C^{-1}B^{\mathsf{T}}, \qquad \gamma^\star \in [0,1]. \tag{5}$$

When the cross-pool block is well estimated the mechanism behaves like a single diversified pool; when it is not, it falls back continuously toward isolated pools, degrading gracefully rather than discretely. No new primitive is required: $A(\gamma^\star)$ is the augmented sub-covariance of Cotton (2024), and the only added ingredient is a running estimate of ρ and n for each pool pair, both available on chain.

6 Discussion

The contribution is deliberately small. Chitra et al. (2025) already reach for the Schur complement and already note its kinship with hierarchical risk parity; what is missing is the damping parameter that turns a binary architectural test into a continuous, self-tuning one. The same move—occupying the interior of a dial that a field has only used at its endpoints—recurs in portfolio construction (HRP versus minimum variance), in spatial statistics (base-shrunk versus full Vecchia conditionals), and in optimization (KFAC's binary block-coupling choices) (Cotton, 2024, 2026). DeFi lending pools are, if anything, the cleanest case for damping, because their estimation problem is the most severe. We leave to future work an empirical comparison of damped versus binary pool architectures on historical PDLP data, and the connection between the dynamic-pricing model and the cost of delta hedging that Chitra et al. (2025) flag as open.

References

T. Chitra, T. Diamandis, N. Sheng, L. Sterle, and K. Yusubov. Perpetual Demand Lending Pools. arXiv:2502.06028, 2025.

P. Cotton. Schur Complementary Portfolios Fix Hierarchical Risk Parity. *Geek Culture* (Medium), 2022.

P. Cotton. Schur Complementary Allocation: A Unification of Hierarchical Risk Parity and Minimum Variance Portfolios. arXiv:2411.05807, 2024.

P. Cotton. Two Sides of Schur Damping: High-Dimensional Pseudo-Likelihoods and Portfolio Allocation. arXiv:2606.14798, 2026.

M. López de Prado. Building Diversified Portfolios that Outperform Out of Sample. *Journal of Portfolio Management*, 42(4), 2016.