SLICR RESEARCH · ETH/USD · BASE · JUNE 2026

# P-Asset Rebalancing Cost Study

Vitalik flags a 2%/yr slippage tax as the biggest risk to option-based index assets. We walked his roll algorithm across 2.23 years of real ETH price and on-chain Uniswap V3 depth. Result: Slicr's TWAP engine today already keeps drag under the bar across every size we tested — and a passive one-sided maker mode could cut it a further ~79%.

Generated 2026-06-02 UTC · Network: Base · Venue: Uniswap V3 WETH/USDC · 19,582 hourly price points · 18 rolls

| 0.80% | under | 0.17% | ~79% | 18 |
|---|---|---|---|---|
| SLICR TWAP $1M | 2%/YR BAR ✓ | MAKER-MODE $1M | MAKER UPSIDE | REAL ROLLS |

| ✓ bar | 0.17% | ~79% | 78 |
|---|---|---|---|
| SLICR TWAP CLEARS 2%/YR | MAKER-MODE $1M DRAG | POTENTIAL MAKER UPSIDE | REAL TICK STATES |

01 — EXECUTIVE SUMMARY

# The slippage tax is real. Active slicing clears the bar; passive making could go further.

On 1 June 2026 Vitalik Buterin proposed building index-tracking assets on top of options rather than debt. The design splits one ETH into a (P, N) pair so the P-asset holder carries stable USD-like exposure with no liquidation risk. The catch he flags himself: holding to maturity re-introduces ETH exposure, so the holder must roll strikes as price moves, and every roll pays slippage. In his words, it is easy to lose 2% per year or more this way, and that is the largest risk by which the whole scheme might become uncompetitive.

We walked his roll algorithm across the real ETH/USD path on Base from January 2024 to March 2026 and priced every roll against real Uniswap V3 tick liquidity. We compared three execution styles, and it matters which one is the product:

**Slicr today is active TWAP slicing** — it splits an order into time-weighted slices and crosses the spread on each one, letting pool depth recover between slices. In this study that is the line labelled **“competent execution.”** It is not a strawman; it is what the product does. Slicr does not post passive liquidity today.

**Passive one-sided making** — posting resting liquidity and waiting for organic flow to fill it over days — is a **different mechanism, and a roadmap extension Slicr has not shipped.** In this study that is the line labelled **“Slicr one-sided.”** Read it as the prize that motivates building a maker mode, not as a current capability.

**Two findings, stated plainly. (1) Slicr's active TWAP slicing today holds $1M annual rebalancing drag at 0.80%/yr and clears Vitalik's 2%/yr bar across every position size below $10M — so the slippage tax does not, on its own, sink the scheme for ordinary holders. (2) A passive one-sided maker mode could cut that drag to ~0.17%/yr in calm markets, about 79% lower, because a patient maker earns the spread instead of paying it. The second number is an opportunity, not a shipped feature.**

The honest comparison throughout is against active slicing — the product as it exists — not against a naive single-venue dump, which no competent operator would do. Every drag figure in this report is the execution cost of the 18 rolls themselves; it does not include the residual quadratic tracking drift that accrues *between* rolls, which is a separate cost the holder also bears.

02 — BACKGROUND: THE PROPOSAL

## Index exposure from options, not debt

One ETH is split into a (P, N) pair with strike S and maturity M. At resolution price x, P receives min(1, S/x) ETH and N receives max(0, 1 - S/x) ETH, so P + N always equals one ETH and no liquidation is ever possible. A P-asset holder's terminal USD payoff is x · min(1, S/x) = min(x, S) — equivalent to a riskless S minus a put struck at S.

To keep stable exposure the holder must not hold to maturity, since that re-introduces ETH exposure. So the holder rolls strikes as the price moves, and each roll is a trade that pays slippage. The open question, quoting the post:

> *“It is very easy to lose 2% per year or more from multiple rounds of slippage, and this is the largest risk by which this whole scheme might become uncompetitive.”*
>
> *“Rebalancing would be more like one-sided market making than like making an instant sell ... users' time preference will almost always be very low.”*

The second quote describes patient, one-sided, low-time-preference execution. That is the mechanism this study finds wins biggest — and it is a maker mode Slicr does not ship today. Slicr's current engine is the active TWAP slicing modelled here as competent execution. The rest of this report measures both: whether active slicing keeps the tax under the bar (it does), and how much further a passive maker mode could push it (a lot, in calm markets).

03 — MODEL & DATA

# All real prices, all real depth

**Input data** (no synthetic prices)

- ETH/USD path: hourly WETH/USD on Base, 2024-01-01 to 2026-03-25 — 19,582 points, despiked for bad oracle prints. Realized volatility 67%/yr.
- Execution depth: real Uniswap V3 WETH/USDC tick snapshots (sqrtPrice, per-tick liquidity_net, fee tier, TVL) across 78 dated states. Slippage walks the actual tick liquidity, not a heuristic constant-product curve.
- Roll schedule: hold P with strike S = X/2.5 and ~45-day maturity; roll into a fresh lower strike when price falls to 1.5×S, and roll at maturity. Walking the real path yields 18 rolls (8.1/yr) over 2.23 years.

## Three execution styles, priced per roll

**Naive (reference only).** Dump the full notional now into the single cheapest WETH/USDC pool. No competent operator does this; shown only as an upper bound.

**Active TWAP slicing — this is Slicr today.** Split the order into 12 slices over ~6 hours, route each slice to the best venue, and let the pool recover between slices. Each slice crosses the spread. This is exactly what Slicr's live engine does, and it is the honest baseline every other number is measured against. Labelled “competent” in the results.

**Passive one-sided making — a roadmap mode, not shipped.** Post one-sided resting liquidity and let organic flow fill it over up to 7 days (low time preference). Cost is adverse selection over the fill window. Any unfilled remainder falls back to active slicing, not a naive dump. Earned market-making fees are not counted as profit — a deliberately conservative choice. Labelled “Slicr one-sided” in the results; read it as the prize, not the product.

**Crowded exit.** The same passive-maker model under synchronized cohort stress: flow capture drops to 1%, patience to 1 day, and adverse selection rises.

**Floor artifact, stated up front.** The crowded-exit number never exceeds active slicing in our table partly by construction: any unfilled passive order falls back to active slicing, which floors the cost at the slicing result. In a real cascade a passive maker can do *worse* than active slicing — you are adversely filled as price drops straight through your resting order. So treat the crowded column as a floored, optimistic bound on the maker mode, not a guarantee it beats slicing under stress.

**Why the drift is quadratic.** P embeds a put, and its gamma makes the hedge adjustment scale with price moves — so cumulative deviation scales with realized variance, the “slow quadratically growing deviation” of the post. Measured annual realized variance here is 0.455 (67%/yr vol). Note this drift is a tracking cost that accrues *between* rolls; the drag figures in this report price the *execution* of the rolls, not this residual.

04 — THE PATH WE WALKED

# Real ETH/USD with every roll event marked

The model is driven by the actual ETH/USD path on Base, not a simulation. Vertical markers show where the roll algorithm fired — 18 rolls over 2.23 years, a mix of price-trigger rolls (when ETH fell through 1.5×S) and scheduled maturity rolls.

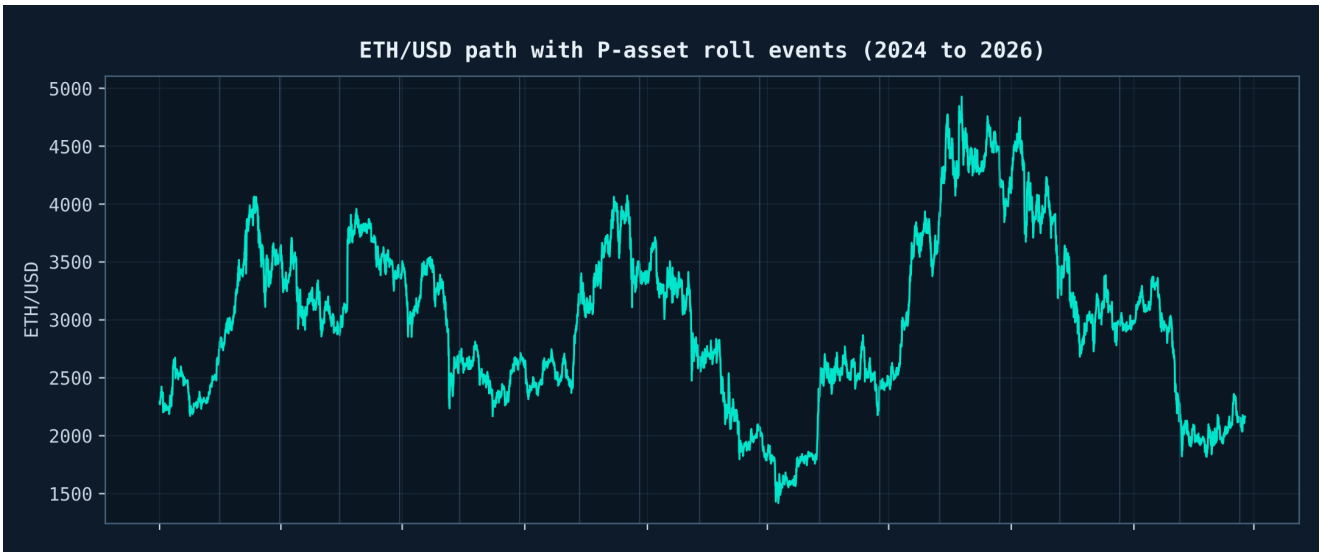


*Real ETH/USD path. Vertical lines mark P-asset roll events.*

Each of those vertical lines is a moment the holder had to trade. The question is simply how much each of those trades costs under each execution model — and the next section prices them.

05 — RESULTS

# Annualized rebalancing drag by position size

Drag is measured against active TWAP slicing — the Slicr engine today. The two right columns are the passive maker mode (calm and crowded); “maker upside” is how much a passive maker could shave off active slicing in a calm market.

| Position | Slicr TWAP (today) %/yr | Maker calm %/yr | Maker crowded %/yr | Calm maker upside |
|---|---|---|---|---|
| $10k | 0.35 | 0.27 | 0.41 | 22% |
| $25k | 0.25 | 0.14 | 0.32 | 44% |
| $50k | 0.29 | 0.10 | 0.34 | 64% |
| $100k | 0.45 | 0.10 | 0.41 | 78% |
| $250k | 0.50 | 0.12 | 0.44 | 76% |
| $500k | 0.63 | 0.14 | 0.54 | 78% |
| $1M | 0.80 | 0.17 | 0.71 | 79% |
| $2.5M | 1.31 | 0.26 | 1.22 | 80% |
| $5M | 1.97 | 0.36 | 1.90 | 82% |
| $10M | 2.90 | 1.07 | 2.85 | 63% |

*Drag versus active TWAP slicing (Slicr today). Green = passive-maker calm and the maker upside; amber = crowded-exit maker drag.*

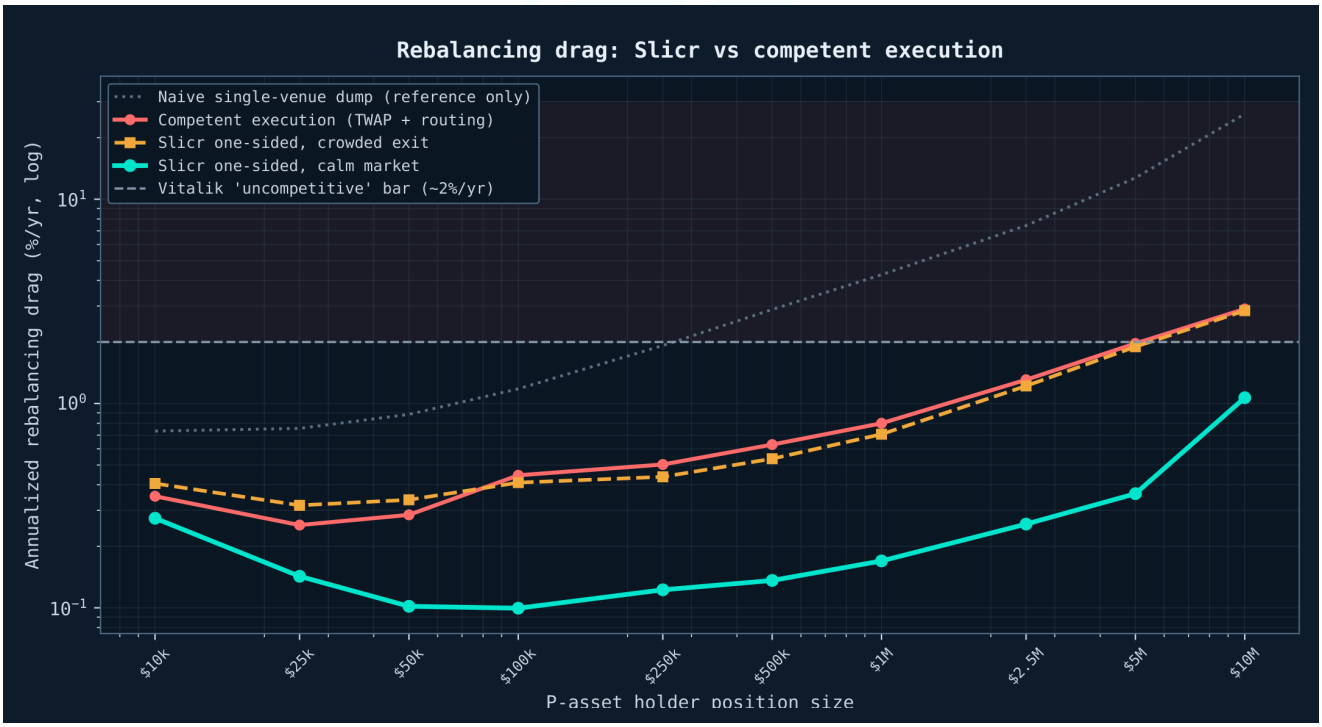


*Drag by holder size on a log scale. “Competent execution” in the chart legend is Slicr's active TWAP slicing today; “Slicr one-sided” is the passive maker mode (roadmap, not shipped).*

## Reading the table

• Slicr's active TWAP slicing today only crosses the 2%/yr bar near $10M. Below that, the engine as it exists already keeps drag under the competitiveness threshold — the slippage tax is real but not, by itself, fatal for ordinary holders.

• A passive maker mode would sit well below active slicing across the whole range in calm markets, because a patient maker earns the spread instead of paying it. That is the ~79% upside — a reason to build the mode, not a number Slicr delivers today.

• In a crowded exit the maker advantage compresses sharply and converges to active slicing at large size. And as noted, the crowded column is floored by the fallback to slicing, so it understates how badly a passive order can fare in a real cascade.

• Below ~$50k, drag is dominated by fixed per-roll gas, not slippage — which is why the small-size curve flattens and ticks up.

**Bottom line: the slippage tax Buterin flags is real, but Slicr's active TWAP slicing already keeps it under the 2%/yr bar for every holder below $10M. The mechanism he points to — patient passive making — could cut drag much further in calm markets, which is a strong case for building a maker mode. But that upside is largest exactly where rebalancing is already cheap, and thinnest in the synchronized stress that matters most. It is an opportunity to pursue, not a free lunch already on the table.**

06 — PER-ROLL COST & FILL RATE

# Where the cost actually lands

Drag is an annualized average, but it is paid roll by roll. The left panel shows the per-roll cost for a $1M holder across all 18 rolls: the first roll is expensive in every model (a large move into thin depth), and the passive-maker calm line sits consistently below both active slicing and the crowded case. The right panel shows that passive-maker calm fills stay near 100% all the way up to $5M, only dropping for a $10M position that organic flow cannot fully absorb inside the patience window.

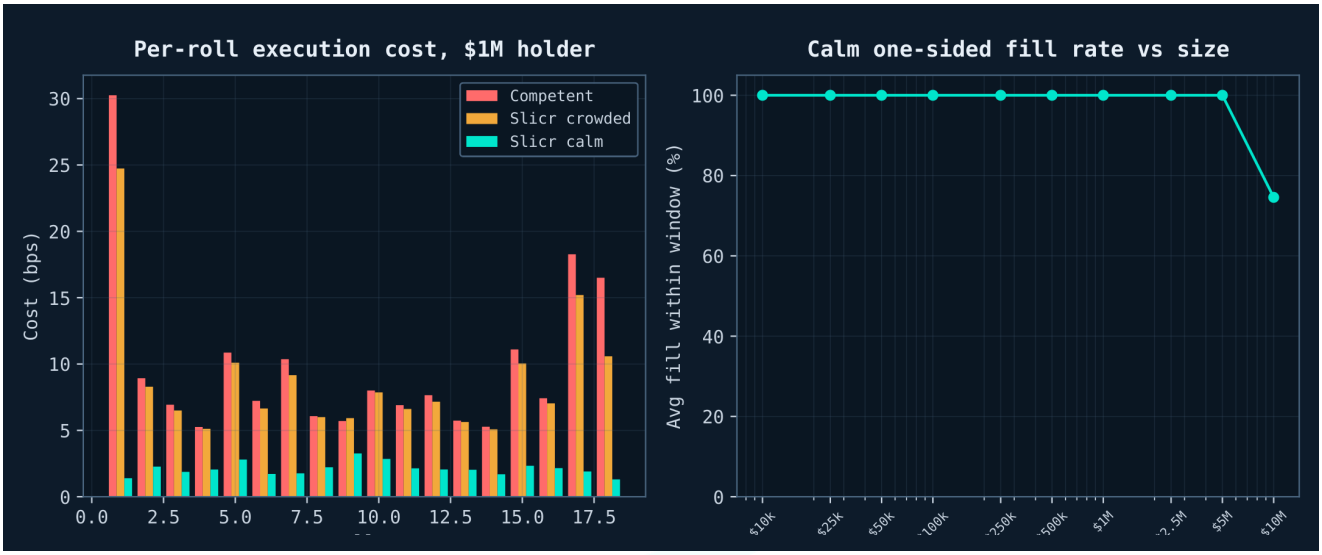


*Left: per-roll execution cost for a $1M holder (legend “competent” = Slicr active slicing today; “calm” = passive maker mode). Right: passive-maker calm fill rate as position size grows.*

The fill-rate cliff at $10M is the same story the drag table tells: at the very top, one-sided making runs out of natural counterflow, the remainder falls back to competent execution, and the edge evaporates.

07 — STRESS & SENSITIVITY

# The crowded-exit problem

The passive maker mode assumes organic counterflow wants the other side of the trade. But every P-holder runs the same roll algorithm. When ETH falls through 1.5×S, the whole cohort wants to sell at once, and the counterflow is one-sided exactly when it is needed. Maturity rolls are staggered across holders and do not crowd; price-trigger rolls do, and that is where a passive maker is thinnest.

Modeled as a crowded exit (1% capture, 1-day patience, raised adverse selection), the maker-mode $1M drag rises from 0.17%/yr to 0.71%/yr, against Slicr's active slicing at 0.80%/yr. At $10M the crowded number (2.85%/yr) is essentially active slicing (2.90%/yr). And recall the floor: that 0.71% sits below slicing partly because unfilled passive orders fall back to slicing. A real cascade, where you are filled into a falling price, can push a passive maker above active slicing — which is precisely why slicing, not passive making, is the safer default under stress.

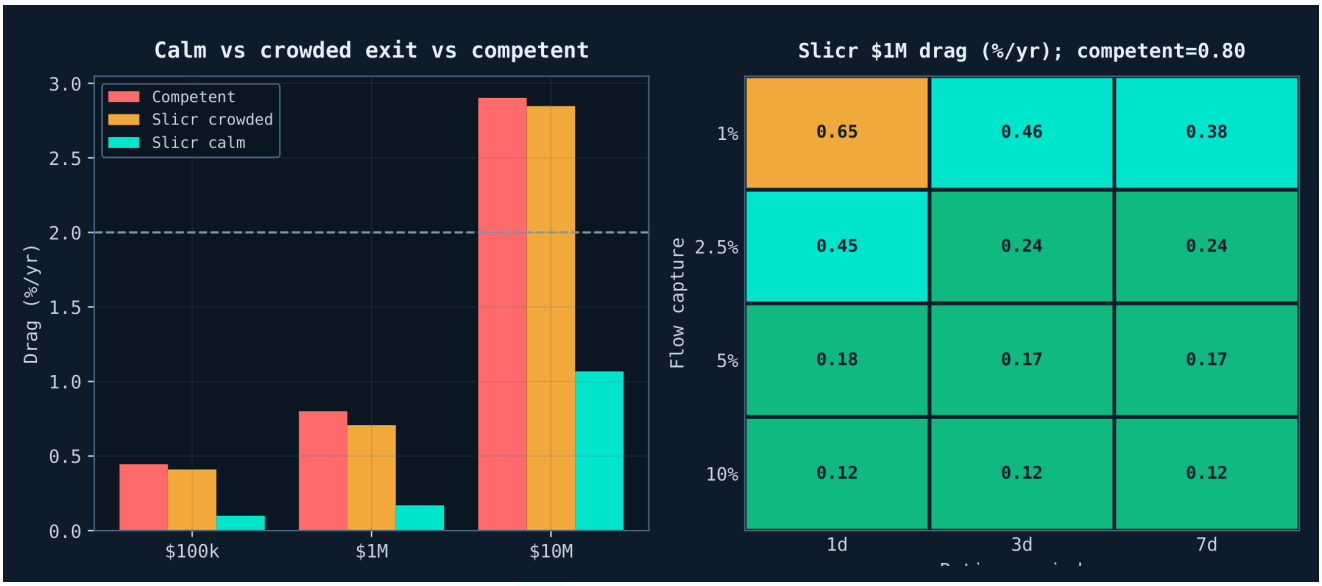


*Left: active slicing vs maker-calm vs maker-crowded. Right: maker-mode $1M drag by flow capture and patience window.*

## Sensitivity grid — passive-maker $1M drag (%/yr)

Two assumptions carry the maker-mode result: how much organic turnover exists, and how much one-sided flow a maker can capture. The grid sweeps capture against patience; Slicr's active slicing at $1M (0.80%/yr) is the reference.

| Capture \ Patience | 1 day | 3 day | 7 day |
|---|---|---|---|
| 1% | 0.65 | 0.46 | 0.38 |
| 2.5% | 0.45 | 0.24 | 0.24 |
| 5% | 0.18 | 0.17 | 0.17 |
| 10% | 0.12 | 0.12 | 0.12 |

*Green cells beat active slicing (0.80%/yr); amber is the closest call. No cell exceeds slicing — subject to the fallback floor noted above.*

**Even at a pessimistic 1% capture and 1-day patience (0.65%/yr), the maker-mode $1M result stays under active slicing (0.80%/yr) — but the margin narrows from large to slim, and the floor flatters it. The maker upside is robust in direction, not in magnitude.**

08 — CAVEATS & SOURCE

# What is and is not modeled

• **Single path, single regime.** This is one realization — one ETH/USD path, n=18 rolls, one volatility regime (67%/yr). The *direction* of every finding is robust, but the *magnitude* is path- and regime-dependent. Higher vol raises roll counts and should widen the gap, but we show that as reasoning, not as a measured result across many paths. A multi-path Monte-Carlo would tighten the confidence intervals; treat the point estimates as illustrative of one history.

• **Depth is real Uniswap V3 WETH/USDC on Base.** Mainnet and aggregated routing are deeper, so the active-slicing baseline is, if anything, a touch pessimistic — which understates how good Slicr's current engine already is, and narrows the maker-mode upside.

• **Active slicing is 12 slices over 6 hours** with best-venue routing and pool recovery — the Slicr engine as it ships. A top-tier external aggregator could do marginally better, narrowing the maker edge further.

• **Passive-maker fills are bounded by assumed organic volume** (turnover ~0.6× TVL; capturing 5% of one-sided flow in calm, 1% in a crowded exit) and charged adverse selection from realized vol. The sensitivity grid shows how much this matters, and the crowded column is floored by the fallback to slicing.

• **Synchronization is the binding risk for a maker mode.** Trigger rolls crowd; maturity rolls do not. The crowded scenario applies the stress to every roll — an upper bound on the crowding penalty in one direction, while the slicing-fallback floor flatters it in the other.

• **Gas is a flat per-roll cost.** It dominates drag below ~$50k and is negligible above, which explains the flat, slightly rising small-size curve.

• **Execution cost only.** Every drag figure prices the execution of the 18 rolls. It excludes the residual quadratic tracking drift that accrues between rolls — a separate cost the holder also pays, and one no execution style removes.

Earned market-making fees are excluded from the Slicr result entirely. Counting them would improve every calm number; leaving them out keeps the comparison conservative and the headline defensible.

## Source

V. Buterin, “Building index-tracking assets on top of options instead of debt”, ethresear.ch, 2026-06-01.
https://ethresear.ch/t/building-index-tracking-assets-on-top-of-options-instead-of-debt/25036

## About Slicr

Slicr is a non-custodial execution engine on Base. It splits large orders into time-weighted slices and actively executes each slice across multiple DEXs, letting pool depth recover between slices, with per-slice price guards enforced in the vault contract to block MEV sandwich attacks. It does not currently provide passive liquidity — it crosses the spread on every slice. In this study, that active slicing is the line that already keeps P-asset rebalancing drag under the 2%/yr bar for holders below $10M.

The larger calm-market saving in this report comes from a passive one-sided maker mode Slicr has not yet built. We are publishing the analysis because it makes a concrete case for that mode: in calm conditions the upside is real and large, while in synchronized stress — the regime that matters most for a P-asset cohort — it compresses and a passive maker can fare worse than slicing. That trade-off is exactly the kind of thing worth resolving in code before claiming it in production. Try the live slicing engine at slicr.xyz/backtest.