Time-Variant AMMs (TVAMMs): Adaptive Liquidity with Passive LPs

Keroshan Pillay[1], Andrew Gonzalez[1], Albert Wildeman[1], Azat Hafizov[2], and Daniel Tang[2]

[1]*Adjusted Finance* , `research@adjusted.finance`
[2]*Pike Finance* , `azat@nuts.finance`; `daniel@nuts.finance`

October 2025

Abstract

We introduce **Time-Variant AMMs (TVAMMs)**, a third DEX category in which the algorithm adapts over time while LPs remain passive. By embedding time in the invariant, pools adjust via standardized governance mechanisms including ramps, envelopes, and circuit breakers. Using **dynamic amplification** (StableSwap's A) as a concrete instance, we demonstrate a TVL-weighted **+392 bps** uplift in mean utility versus the constant-product invariant baseline and the best static settings, with stable 1% Shortfall VaR (SVaR) and no change in ΔIL. We formalize the category, outline guardrails, and invite curators/operators and asset issuers to participate in the TVAMM roadmap. In this platform, StableSwap with dynamic amplification $A(t)$ is the first supported algorithm and benchmarked policy, not the category itself.

Contents

1 Executive Summary . . . 2

2 Background and Related Work . . . 4

3 How Time-Variant AMMs Work (High-Level) . . . 5

3.1 Backtest Scope & Limitations . . . 5

4 Backtest Results . . . 6

5 Mechanism and Market-Structuring Effects . . . 8

6 Policy Design for TVAMMs . . . 9

7 Risk Assessment . . . 10

8 Implementation Considerations . . . 12

9 Open Questions and Future Work . . . 12

10 Conclusion . . . 13

A Utility Function & Risk Metrics . . . 14

B Notation and Symbols . . . 15

C Derivation of the Curvature Penalty Term . . . 15

D Reproducibility and Benchmarks (Backtest) . . . 16

1 Executive Summary

We introduce a third DEX category—*AMM proactive / LP passive*—where the protocol manages time-aware parameters under standardized safety mechanisms (*TVAMM rails*) while LPs remain passive. TVAMMs embed time in the invariant so pools can adapt via governed ramps, envelopes, and circuit breakers. StableSwap with dynamic amplification $A(t)$ serves as the first supported algorithm within this platform—a concrete policy instance, not the category itself.

CLMMs improved capital efficiency but fragmented liquidity and demand continual rebalancing from LPs. TVAMMs preserve efficiency gains while restoring the passive-LP paradigm: cohesive liquidity with predictable risk and standardized governance safeguards.

Our Jan–Apr 2025 backtest measures a TVL-weighted **+392 bps** uplift in mean utility versus the constant-product invariant baseline and best static A configurations, with stable SVaR (1%) and unchanged ΔIL. Using dynamic amplification (StableSwap's A) as the concrete policy instance, this single-snapshot statistical analysis demonstrates category viability. Path-dependent ramps and live governance are reserved for future phases. The AMM Quadrant (Figure 1) positions TVAMMs as the third category of DEX and anchors the terminology used throughout.

TVAMMs increase the set of economically viable larger trades by expanding the near-peg linear region, raising fee potential per TVL. They restore passive LP participation without continual rebalancing and provide extensible infrastructure that generalizes beyond stable pairs. The category enables protocol-level optimization while keeping liquidity providers passive—and better paid.

Definition: TVAMM A Time-Variant AMM has an invariant $f(\text{reserves}, t) = k$ with time-aware policy control (e.g., $A(t)$) and governance guardrails (ramps, envelopes, circuit breakers). Category: *AMM proactive / LP passive*. *Platform note: StableSwap with dynamic amplification $A(t)$ is the first supported algorithm/policy, not the category itself.*

Glossary: Quick Reference

- **TVAMM**: Category where the protocol manages time-aware parameters under governance; *AMM proactive / LP passive*.
- **TVAMM rails**: Standardized category-level safety mechanisms (ramps, envelopes, circuit breakers, jitter) that generalize across all time-variant policies.
- **Curator/Operator**: Team that designs and executes the parameter-adjustment policy.
- **Ramp**: Bounded change over time; **Envelope**: Min/max bounds; **Dwell time**: Minimum on-chain duration between changes.
- **Circuit breaker**: Halt/revert conditions (e.g., oracle anomaly, extreme de-peg).
- **Near-peg flat (linear) region**: Low-curvature region with minimal slippage for small-to-moderate trades.
- **Dynamic A**: Time-aware adjustment of StableSwap's amplification parameter; *first supported algorithm/policy under TVAMM*.
- **SVaR (1%)**: Shortfall Value at Risk at the 1% tail of the utility distribution.
- **ΔIL**: Relative impermanent loss versus the *constant-product invariant baseline*.
- **Trade-size bands**: Discrete bands of trade sizes used for analysis and policy.
- **Oracle delay**: Latency window used in volatility-loss modeling and safety logic.
- **Decision gate**: Multi-criteria gate to promote research policies to production.
- **bps**: Basis points (1 bps = 0.01%).

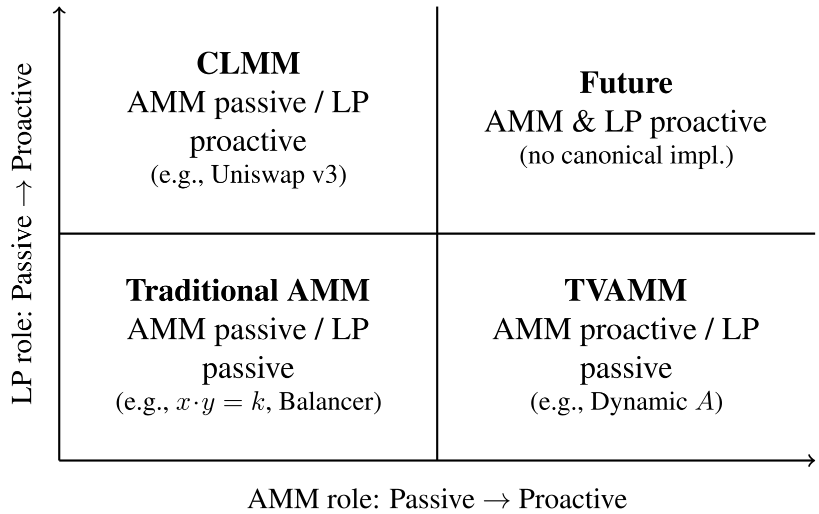

Figure 1: The AMM Quadrant: a categorical (non-spectral) taxonomy that anchors the paper.

Backtested, simulated results; not financial advice.

2 Background and Related Work

CLMMs trade cohesive liquidity for LP control and require ongoing maintenance; liquidity often fragments or drifts out of range. Passive LPs seek cohesive liquidity and predictable risk without 24/7 management.

CFMMs and invariants. Constant-Function Market Makers (CFMMs) are defined by an invariant $f(\text{reserves}) = k$; prices follow from the invariant's gradients. The *constant-product invariant baseline* ($x \cdot y = k$) and StableSwap are canonical CFMMs referenced throughout [1, 2].

StableSwap and amplification. StableSwap introduces an amplification parameter A that concentrates liquidity near the peg, expanding a low-curvature *linear regime* where small-to-moderate trades incur near-zero slippage [1]. Larger A flattens curvature near equilibrium but can increase exposure once price exits the flat region.

AMM Quadrant: a categorical map. We classify DEX designs along two axes—AMM role (passive vs. proactive parameter management) and LP role (passive vs. proactive liquidity management)—yielding four distinct categories (Figure 1):

1. **AMM passive / LP passive** (Traditional AMM): parameters are effectively static; e.g., $x \cdot y = k$, Balancer weighted pools.
2. **AMM passive / LP proactive** (CLMM): AMM parameters are static while LPs actively manage ranges; e.g., Uniswap v3.

3. **AMM proactive / LP passive** (TVAMM): the protocol dynamically manages time-aware parameters under standardized governance guardrails—*TVAMM rails*—while LPs remain passive. Dynamic A exemplifies this category.
4. **AMM proactive / LP proactive**: a future paradigm without canonical implementations.

Positioning. This whitepaper establishes Quadrant 3 as a distinct category. TVAMMs employ standardized safety mechanisms—ramps, envelopes, circuit breakers, and timing controls—that constitute reusable *TVAMM rails*. StableSwap with dynamic amplification $A(t)$ serves as the first concrete policy implementation operating within this framework, demonstrating the category's viability. CLMM comparators are intentionally deprioritized to isolate protocol-level optimization within the AMM-proactive paradigm.

3 How Time-Variant AMMs Work (High-Level)

Definitions. A **Time-Variant AMM (TVAMM)** has an invariant $f(\text{reserves}, t) = k$ with explicit time dependence via policy-controlled parameters (e.g., $A(t)$). All TVAMMs operate on **TVAMM rails**—standardized category-level safety mechanisms including ramps, envelopes, circuit breakers, and timing controls that generalize across any time-variant policy. *Dynamic* A denotes a specific instantiation where the amplification parameter updates according to a policy subject to these governance and safety constraints.

Intuition first. Think *governed ramps on rails* that smoothly update parameters in response to market state—*without* asking LPs to chase price. Keep liquidity cohesive; let the protocol do the work. Figure 5 offers an intuitive view of how amplification flattens curvature near the peg and expands the linear region in practice.

3.1 Backtest Scope & Limitations

The backtest is *single-snapshot* and *strictly statistical*. It benchmarks: (i) the *constant-product invariant baseline*, (ii) a grid of static A values, and (iii) a selector by trade-size band (policy chooses A by band). Price paths follow GBM; oracle delay Δt is incorporated via a volatility penalty. **Paths/config:** 50k; **total configs:** 576; **random seeds:** fixed per experiment (see repo). The 3Pool is modeled as a symmetric two-asset pool near equilibrium for analytical consistency. Out of scope: multi-pool interactions, adversarial manipulation, sequential transaction effects, and on-chain ramp mechanics. Next steps address path dependence and governance-constrained ramps.

Scope of evidence. Dynamic $A(t)$ is used as a *case study within TVAMM*; results do not imply that TVAMM is restricted to StableSwap.

4 Backtest Results

Backtest Evidence Snapshot. Four high-TVL pools (Jan 1–Apr 30, 2025); 50k paths/config; 12 A values; 12 trade-size bands; 576 configs. Dynamic A achieves a +392 bps TVL-weighted uplift versus the constant-product invariant baseline and the best static A; SVaR (1%) remains stable; ΔIL is unchanged. The 3Pool is modeled as a symmetric two-asset pool near equilibrium. The production solver (100 iterations) is validated to machine precision; benchmarks appear in the Appendix.

Top-line results. Utility analysis (Figures 2a–2b) reveals band-dependent optima and a widening linear region as A increases. We benchmark dynamic $A(t)$ within the TVAMM category against the constant-product invariant baseline and the best static A. The Decision Gate in Figure 3 shows a TVL-weighted +392 bps improvement under the dynamic policy versus baselines. Figure 4 indicates that 1% SVaR is stable across A. (For the precise utility function definition and risk metrics, see Appendix A.)

Translating Utility to LP Revenue: A Framework

Utility serves as a research proxy for economic value, but liquidity providers care about fee/TVL. The mechanism is straightforward: an expanded near-peg linear region attracts larger economically viable trades, directly increasing fee income per unit of TVL. The realized uplift depends on fee rate f, trade composition, and execution dynamics.

Revenue mapping framework. Let baseline revenue be r_f (fee/TVL per horizon H, e.g., a 30-day trailing rate) and the policy's utility uplift be $+\Delta u$ (e.g., 392 bps = 3.92%). The revenue uplift can be modeled as

$$\Delta r_f \approx \alpha \, \Delta u \, r_f, \quad \alpha \in (0, 1),$$

where α captures the conversion efficiency from utility gains to realized fees, accounting for trade mix, execution quality, and liquidity migration effects. As an *illustrative example only*, if $\alpha = 0.5$, a 392 bps utility uplift translates to a first-order revenue increase of $\approx 1.96\% \times r_f$. The actual α is implementation-specific and must be calibrated by the operator per pool (using realized fee/TVL and trade-mix data) and approved through governance.

Pool	**Mean TVL (m USD)**	**Trades**	**Δ vs Best Static (%)**
3Pool	182.0	45,097	3.8%
USDC/DAI	57.1	25,671	5.4%
USDC/ETH	152.0	707,969	1.1%
stETH/ETH	127.1	30,996	6.7%
TVL-weighted	—	—	**3.92%**

Table 1: Per-pool improvement of dynamic A versus the best static A, with context (mean TVL and trade count from the Phase-1 report). TVL-weighted aggregate uplift is +3.92% (392 bps). Bootstrap validation uses 2,000 resamples ($p < 0.01$).

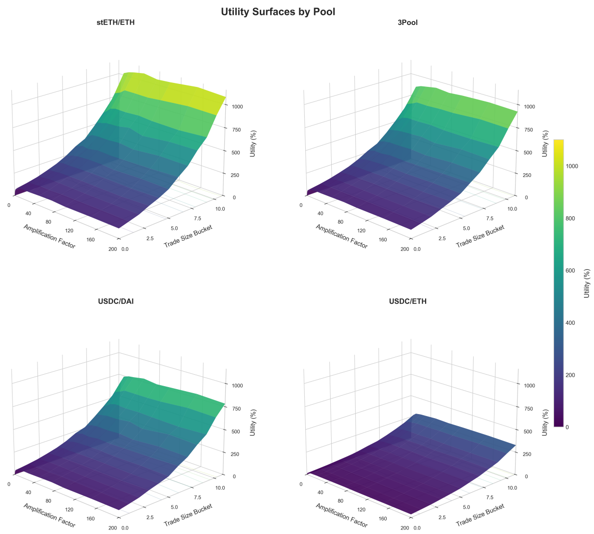

(a) Utility surfaces by amplification and trade-size bands.

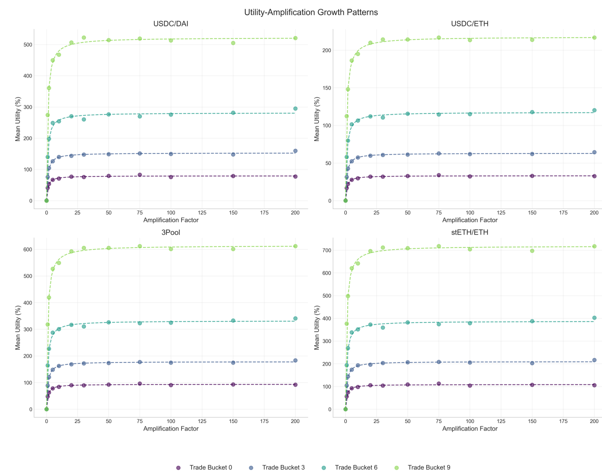

(b) Mean-utility growth by pool and trade size.

Figure 2: Utility analysis across amplification levels and trade-size bands.

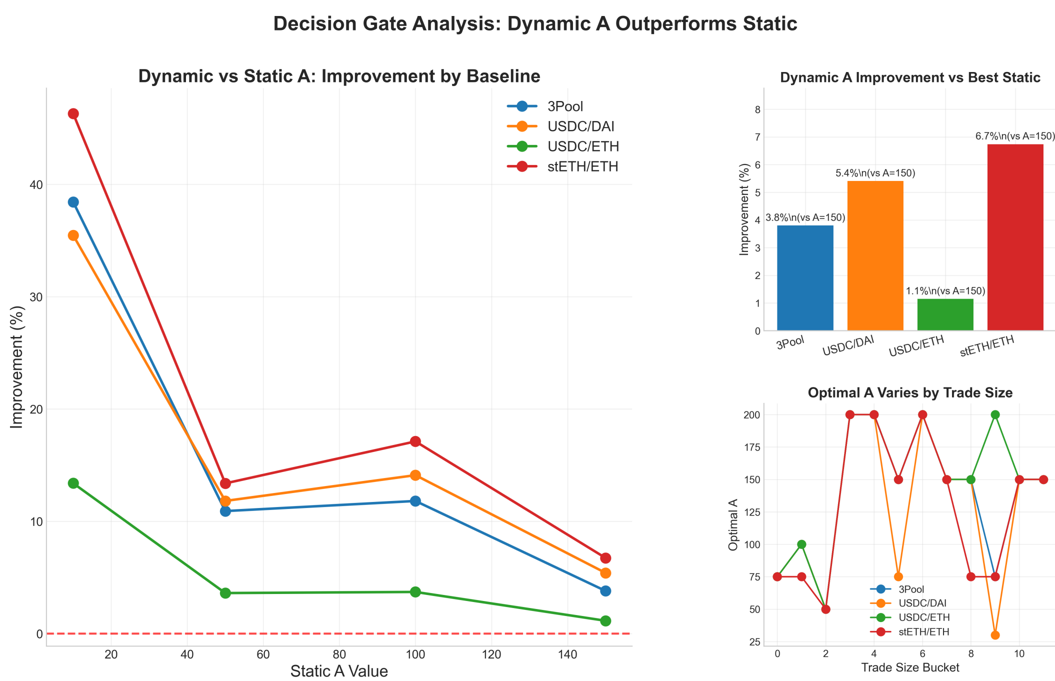

Figure 3: Decision Gate: dynamic A delivers a TVL-weighted +392 bps mean-utility uplift versus the *constant-product invariant baseline* and the best static A across all pools. 3Pool modeled as symmetric two-asset for analytical consistency.

Note: 95% bootstrap confidence intervals (CIs) for per-pool uplifts are provided in Appendix D.

Modeling note. Trade sizes follow a log-normal body with Generalized Pareto tails (95th-percentile POT), and utility is TVL-normalized.

- **Body–tail fit:** log-normal body + GP tails (POT@95th); utility TVL-normalized.

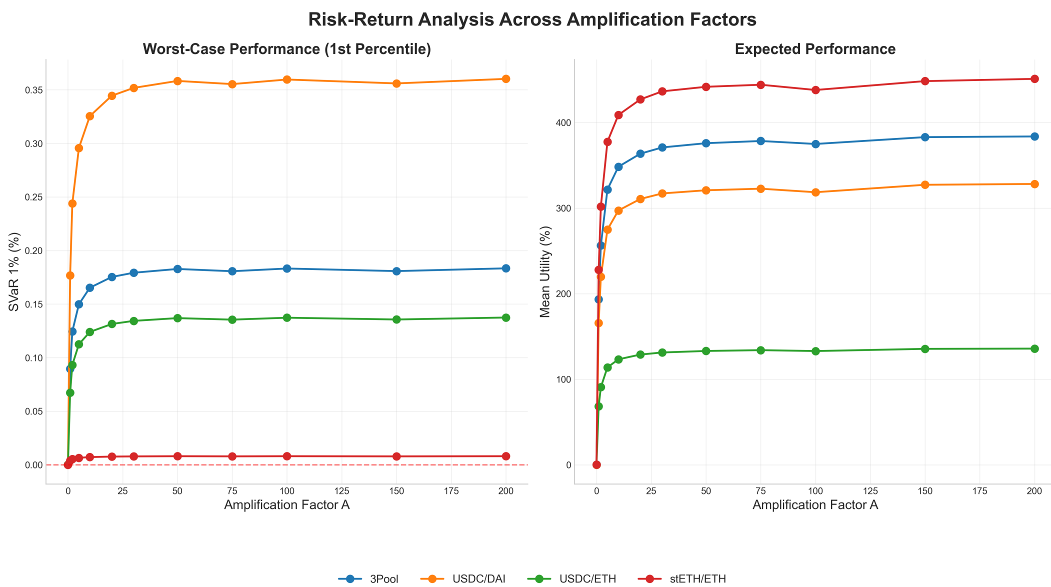

Figure 4: Risk–return profile: SVaR (1%) remains stable across amplification while mean utility increases, indicating no deterioration in tail risk.

5 Mechanism and Market-Structuring Effects

Scope. This section analyzes category-level TVAMM mechanics. Throughout, we use StableSwap with dynamic amplification $A(t)$ as a *policy instance* to illustrate the effects; the mechanics themselves generalize to any time-variant invariant $f(\cdot, t)$ operated on TVAMM rails.

Linear-regime expansion. Increasing A flattens curvature near the peg, expanding a low-slippage linear regime. This reshapes the feasible set of trade sizes that are economically viable given fees and price impact. As Figure 5 illustrates, higher A creates a wider flat region, and observed trade-size density concentrates where slippage is tolerable. This linear-region expansion is a category effect of TVAMM; $A(t)$ is used here as the illustrative policy.

Order-flow composition and fees. Because fees scale with notional, attracting larger trades increases fee income per unit of TVL. The utility analysis (Figure 2) shows diminishing returns to A for small trades but continued gains for large trades; a dynamic policy selects higher A for larger buckets and lower A for smaller buckets, improving average utility without increasing tail risks (Figure 4).

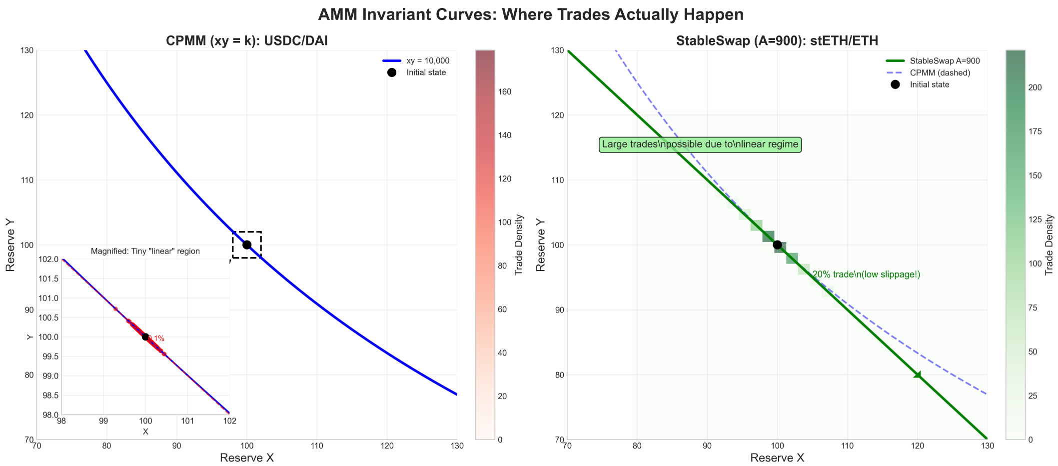

Figure 5: StableSwap invariant curves *(as a TVAMM policy instance)* with trade-size density overlay showing how increasing A enlarges the near-peg linear region, enabling larger economically viable trades and higher fee capture per TVL.

Selector logic. For each trade-size band b, the policy selects $A^*(b) = \arg\max_A \mathbb{E}[U(A, b)]$ on the measured utility surfaces, subject to TVAMM rails (envelopes $[A_{\min}, A_{\max}]$ and governed ramps with minimum dwell). This band-wise selection is used for the dynamic-A case study within the category.

6 Policy Design for TVAMMs

Objectives. Maximize TVL-weighted mean utility subject to SVaR (1%) and governance constraints. Near-term extensions include fee-centric variants that explicitly target fee/TVL and *dynamic swap-fee* policies jointly optimized with A.

TVAMM rails: standardized safety mechanisms. All TVAMM policies operate on **TVAMM rails**—category-level safety features that generalize across any time-variant invariant $f(\cdot, t)$. These include governed ramps (bounded slope, minimum dwell), parameter envelopes $[A_{\min}, A_{\max}]$, circuit breakers (oracle anomaly, extreme de-peg), and timing controls such as jitter (see [3]). Dynamic A leverages these rails; future policies inherit the same robust framework.

Tail-risk budgets. Maintain SVaR (1%) below a governance threshold, with fast fallback to A^{safe} if breached.

Incentives and Keeper Cost

This subsection provides an algebraic sketch to reason about incentives and operating costs; quantitative calibration is left to the operator using pool-specific data and governance review.

Incremental revenue lens. Let baseline revenue be R_{baseline} (fee/TVL per horizon H, e.g., 30-day) and the policy's utility uplift be $+\Delta u$. A compact proportional sketch is

$$\Delta R \approx \alpha \, \Delta u \, R_{\text{baseline}}, \quad \alpha \in (0, 1),$$

where α captures mapping from utility to realized fees (mix, execution, migration).

Costs (components).

- On-chain transactions (gas)
- Off-chain compute/monitoring
- Curator time

Denote total operating cost by C_{ops}.

Break-even (governance cue). Consider promotion when $\Delta R \geq C_{\text{ops}}$ (both measured over the same horizon H) under governance, with conservative defaults.

Keeper rewards. Options include (i) a share of ΔR, or (ii) a fixed stipend tied to service-level objectives (SLOs); any scheme should be subject to governance approval and on-chain constraints.

Note. This is an order-of-magnitude sketch, not a precise mapping; operators should calibrate with observed fee/TVL and trade-mix data.

Decision gate. A straightforward gate (Figure 3) governs promotion to production: (i) mean-utility lift versus baselines, (ii) SVaR stability, (iii) operational feasibility, and (iv) governance readiness.

Curation as art + science. Operators bring distinct styles under shared TVAMM rails. Start conservatively—low-frequency updates and narrow envelopes—then increase cadence as signal quality and monitoring mature. The rails provide safety; the policy provides edge.

Operator Policy Checklist

Objective	Maximize TVL-weighted mean utility subject to Shortfall VaR (SVaR, 1%) $\leq$ threshold
Inputs	Trade-size mix; realized slippage vs model; volatility band; imbalance
Controls	$A_{\min}, A_{\max}$; ramp slope; min dwell; circuit-breakers
Cadence	Curator preference based on market conditions and monitoring maturity
Fallbacks	Revert to A^{safe} on breaker trip; freeze updates on oracle anomaly
Monitoring	Fee/TVL; trade composition; SVaR proxy (rolling 1% quantile); model residuals

7 Risk Assessment

Impermanent loss and tails. Under the calibrated penalties and dynamic selection in this backtest, ΔIL remains unchanged and 1% SVaR is stable across A (Figure 4).

Adversarial considerations and TVAMM rails. The standardized safety mechanisms—*TVAMM rails*—provide systematic defenses against adversarial vectors. These are category-level features, not policy-specific patches:

- **Oracle games near ramps:** anomaly detection and hysteresis; freeze on diverging feeds.
- **Predictable-A sandwiching:** randomized dwell/jitter; bounded envelopes.
- **Liquidity migration:** cross-pool monitoring; throttle cadence on sharp TVL shifts.

Vector	**Scenario**	**Mitigations (TVAMM rails)**
Predictable-A sandwiching	Windows during governed ramps where future A is anticipated	Randomized dwell/jitter; bounded envelopes; capped ramp slopes; delayed publication
Oracle/ramp manipulation	Stale/divergent feeds, de-peg, adversarial sTapios	Redundant oracles with hysteresis; circuit breakers $\to A^{\text{safe}}$
Curator misbehavior	Error or malice in parameter updates	Multisig; attestations; timelocks; on-chain constraints; rollback

Table 2: Threat Model Snapshot (non-exhaustive): selected vectors and mitigations tied to TVAMM rails.

Predictability in ramp timing and level can create MEV windows (*predictable-A sandwiching*). TVAMM mitigates this with randomized dwell/jitter, bounded envelopes, capped slopes, and delayed publication, reducing anticipatory edge.
Oracle/ramp manipulation is addressed via redundant oracles with hysteresis and circuit breakers to A^{safe}. Curator misbehavior is contained through multisig, attestations, timelocks, and on-chain constraints.

Category-level safety: TVAMM rails as systematic risk management. The mitigations detailed above—ramps, envelopes, circuit breakers, jitter, and timing controls—constitute **TVAMM rails**: standardized, category-level safety features that generalize to any time-variant invariant $f(\cdot, t)$, not merely dynamic A. This architecture ensures that all future TVAMM policies inherit robust defenses against oracle manipulation, MEV exploitation, and liquidity shocks. The rails represent a systematic risk-management philosophy: policy innovation occurs within a pre-audited safety framework.

Predictable-A sandwiching. Ramp predictability can create MEV windows. Cadence control, jitter, bounded envelopes, and capped slopes—all TVAMM rails—reduce the exploitable surface.

Oracle triggers. Breakers respond to:

- Feed divergence beyond tolerance or abnormal latency.
- De-peg or volatility sTapio relative to rolling bands.

- Freeze to A^{safe} on anomaly; resume after reconciliation.

These triggers align with the *Monitoring* paragraph and circuit-breaker rails. Beyond algorithmic risks, robust operational security is essential, including safeguards against deployment-time threats (e.g., proxy hijacking) and strong key-management hygiene.

Pause triggers. Oracle inconsistency, extreme de-pegs, or unexplained SVaR drift should pause updates and revert to A^{safe} while operators investigate. Ramp time affects both efficacy and safety; conservative dwell times and slopes reduce exploitability (see [3]).

Monitoring. Track (i) fee/TVL, (ii) trade-size composition, (iii) realized slippage versus model, (iv) SVaR proxies (e.g., rolling 1% utility quantiles), and (v) deviation triggers tied to circuit breakers. These monitoring controls integrate with TVAMM rails and apply to any time-variant policy.

8 Implementation Considerations

Patterns. Compute-and-commit policy with attestations; contracts enforce TVAMM rails (ramps, envelopes, circuit breakers); redundant oracles with hysteresis; batch updates at a chosen cadence. Detailed implementation playbooks will be provided as a separate Tapio implementation memo; this document maintains a platform focus. The architecture generalizes to any time-aware invariant $f(\cdot, t)$, with rails enforced at the contract layer.

9 Open Questions and Future Work

- **Dynamic swap fees:** fee-centric objectives and joint A–fee policies (Tapio will explore this next).
- **Next steps: sequential dynamics:** decision rules over transaction sequences with governance-constrained ramps.
- **TVAMM generalization:** beyond dynamic A to fully time-variant invariants $f(\cdot, t)$. StableSwap $A(t)$ is one supported policy today; additional $f(\cdot, t)$ policies are future work.
- **Cross-pool and game theory:** liquidity migration, competitive responses, and equilibrium effects.
- **Risk metrics:** higher-order moments, Conditional VaR, and drawdown-aware objectives.
- **Adversarial analysis:** formal threat models for oracle and ramp manipulation.
- **Operator analyses:** LP revenue uplift calibration and curator/keeper incentives (as referenced in §4 and §6) should be developed by operators for their pools and reviewed through governance.

10 Conclusion

Dynamic amplification reframes A from a static tuning knob into a market-design instrument. This backtest demonstrates a TVL-weighted +392 bps uplift in mean utility versus strong baselines, with stable SVaR (1%) and no change in ΔIL. The observed utility uplift follows category mechanics—protocol-managed expansion of the near-peg linear region—illustrated here via $A(t)$, which attracts larger economically viable trades and raises fee/TVL.

For launch, we recommend a conservative dynamic-A policy operating on TVAMM rails: explicit ramps, envelopes, circuit breakers, and continuous monitoring. Next steps extend to path-dependent policies and dynamic fees. In this platform, StableSwap with dynamic amplification $A(t)$ is the first supported algorithm and benchmarked policy, not the category itself. The broader TVAMM agenda places the protocol in the AMM-proactive / LP-passive quadrant, enabling protocol-level optimization while keeping LPs passive—and better paid.

TVAMMs represent the next step in DEX evolution—delivering protocol-level optimization while returning the liquidity provider to their rightful role: passive, protected, and profitable.

Disclaimer. This is a research/engineering discussion, not financial advice. Results depend on market conditions and implementation choices.

A Utility Function & Risk Metrics

To rigorously evaluate the performance of different amplification strategies, we developed a comprehensive utility function that captures the multifaceted economic impacts of a trade. The expected utility U of a trade is a function of the amplification parameter A, the trade size Δx, and the asset volatility σ:

$$U(A, \Delta x, \sigma) = \underbrace{(p_{\text{CPMM}} - p_{\text{exec}})}_{\text{Slippage Gain}} - \underbrace{\lambda \left| \frac{\partial^2 p}{\partial \tau^2} \right|}_{\text{Curvature Loss}} - \underbrace{\gamma \, \sigma^2 \, \Delta t}_{\text{Volatility Loss}} + \underbrace{\frac{f \, \Delta x}{\text{TVL}}}_{\text{Fee Income}} \tag{1}$$

Component breakdown.

Slippage Gain: Measures the price improvement a trader receives from the StableSwap invariant compared to a baseline Constant Product Market Maker (CPMM). Higher values indicate better execution quality.

Curvature Loss: Penalty for non-linear price impact. High curvature means execution price accelerates away from market price rapidly as trade size increases, signaling higher impermanent loss risk. See Appendix C for derivation.

Volatility Loss: Captures the cost of impermanent loss risk, proportional to asset price variance (σ^2). Reflects fundamental risk LPs undertake.

Fee Income: Direct revenue generated for LPs from a trade. Normalized by TVL for consistent cross-pool comparison.

Risk metrics.

- **Shortfall Value at Risk (SVaR, 1%):** The 1st percentile of the utility distribution, representing expected outcome in the worst 1% of scenarios. Measures tail risk.
- **Impermanent Loss Difference (ΔIL):** Relative IL between StableSwap and the constant-product invariant baseline, measured across price movements.

Calibrated parameters. $\lambda = 50\,\text{bp}^{-1}$ (curvature penalty coefficient); $\gamma = 0.002$ (volatility penalty coefficient); calibrated from historical pool performance to balance slippage gains against risk exposure.

B Notation and Symbols

Symbol	Definition	Economic meaning
A	StableSwap amplification parameter	Liquidity concentration knob; enlarges linear region
D	StableSwap invariant level	Determines reserve states for a given A
Δx	Trade size (notional)	Order-flow magnitude
σ	Annualized volatility	Market risk proxy
λ	Curvature penalty (λ)	Penalizes accelerating price impact (curvature)
γ	Volatility penalty (γ)	Penalizes IL risk via variance
f	Swap fee rate	Direct LP revenue per trade
TVL	Total value locked	Pool scale; normalizes fee income
p_{xyk}	Execution price under $x \cdot y = k$	Baseline comparator for slippage gain
p_{exec}	Achieved execution price	Realized price in the target invariant
τ	Trade progress ($0 \to 1$)	Parameterization along a single swap path
Δt	Oracle delay window	Evaluation horizon for the volatility term
A^{safe}	Safe fallback amplification	Governance-defined parameter used under circuit-breaker/freeze conditions
C_{ops}	Operating cost	On-chain gas + off-chain compute/monitoring + curator time (per stated horizon)

Table 3: Notation and economic interpretation used throughout.

C Derivation of the Curvature Penalty Term

Let $p(x)$ denote the marginal price (the negative slope of the invariant curve). Define *normalized trade progress* $\tau \in [0, 1]$ along a trade of size Δx from reserve x_0, so $x(\tau) = x_0 + \tau \Delta x$. By the chain rule,

$$\frac{\partial p}{\partial \tau} = \frac{\partial p}{\partial x} \cdot \frac{\partial x}{\partial \tau} = p'(x(\tau)) \cdot \Delta x, \quad \frac{\partial^2 p}{\partial \tau^2} = p''(x(\tau)) \cdot (\Delta x)^2.$$

Thus the penalty term is

$$\lambda \left| \frac{\partial^2 p}{\partial \tau^2} \right| = \lambda \, |\kappa(x(\tau))| \, (\Delta x)^2,$$

where $\kappa(x) = p''(x)$ is used as a *proxy* for geometric curvature. Strictly, arc-length normalized curvature is $\kappa_g = \dfrac{|p''(x)|}{\left(1 + (p'(x))^2\right)^{3/2}}$. Near the peg, $|p'(x)| \ll 1$, so $\kappa_g \approx |p''(x)|$ and the proxy is accurate; away from the peg this overstates curvature and should be treated as an approximation. In a two-asset StableSwap near equilibrium with large A, curvature scales down approximately as $\kappa \propto 1/A$, explaining why higher A suppresses curvature in the linear regime.

D Reproducibility and Benchmarks (Backtest)

Experiments span 12 amplification settings $\times$ 12 trade-size bands $\times$ 4 pools, with 50,000 Monte-Carlo paths per configuration. Price paths follow GBM; oracle delay Δt is applied in the volatility penalty. **Seeds:** PRNGs are fixed for reproducibility: JAX uses `PRNGKey(42)` and NumPy uses `seed(42)` where applicable; the pipeline passes a single `random_seed=42` per experiment. **Solver:** the StableSwap state is solved via a Newton–Raphson routine with 100 iterations and absolute-error tolerance $< 10^{-12}$. Bootstrap standard errors use 2,000 resamples; decision-gate metrics report TVL-weighted averages. The full simulation is available at https://github.com/adjusted-finance/tapio-tvamm, and a Streamlit Dashboard showing cached results can be found at https://tapio-tvamm.streamlit.app

References

[1] Michael Egorov. StableSwap – efficient mechanism for stablecoin liquidity. Curve Finance Whitepaper, 2020. URL: `https://curve.fi/files/stableswap-paper.pdf`.

[2] Guillermo Angeris and Tarun Chitra. Improved price oracles: constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (AFT '20)*, 2020.

[3] Curve Team. Parameter Ramping Considerations. Blog post, 2023. URL: `https://blog.curvemonitor.com/posts/parameter-ramping/` (accessed 2025-09-14).

[4] Paradigm Research. TWAMM: Time-Weighted Automated Market Makers. Blog post, 2021. URL: `https://www.paradigm.xyz/2021/07/twamm`.