

BasisOS: A Managed Basis Trading Strategy for Yield Optimization in Exotic Markets

Boring Wozniak
xboringwozniak@logarithm.fi

Numa Pumpilius
numa@logarithm.fi

February 18, 2025

Abstract

We present a comprehensive study of a managed basis trading strategy implemented via the BasisOS protocol. This strategy captures funding fees using a dual-position approach that simultaneously maintains a long exposure in the spot market and a short exposure on perpetual decentralized exchanges. In contrast to centralized mechanisms (e.g., Ethena[1]), which emphasize total value locked (TVL) and rely on custodian-based systems, the BasisOS vault framework focuses on yield optimization in exotic markets (e.g., Perp DEXs, memecoins) with an emphasis on transparent execution. In this article, we detail the system architecture, introduce rebalancing algorithms, discuss risk management, outline vault capacity estimation methods, and report extensive simulation and backtesting results. Our analysis demonstrates that precise rebalancing combined with robust risk controls is critical for competitive performance in volatile markets.

1 Introduction

Basis trading strategies exploit funding rate imbalances by simultaneously maintaining a spot position and a corresponding hedge (short) position in the perpetual futures market. The primary objective is to capture periodic funding fee premiums while maintaining a target leverage L_{target} . This target leverage governs the vault's utilization and is confined between a minimum leverage L_{min} (ensuring optimal funding collection) and a maximum leverage L_{max} (limiting risk exposure).

$$L = \frac{|PositionSize|}{MarginBalance}, \quad L_{\text{target}} \in (L_{\text{min}}, L_{\text{max}})$$

Market Overview

Exotic markets typically offer elevated funding rates compared to blue-chip markets, albeit with

increased volatility and liquidity challenges. We quantify these funding rates using two indices (see [4]):

- **Average Funding Rate Index (AFRI):**

$$AFRI = \frac{1}{N} \sum_{i=1}^N r_i,$$

where r_i is the annualized daily funding rate of the i -th market and N is the total number of markets.

- **Weighted Funding Rate Index (WFRI):**

$$WFRI = \frac{\sum_{i=1}^N (r_i \cdot OI_i)}{\sum_{i=1}^N OI_i},$$

where OI_i denotes the daily average open interest of the i -th market.

We further define the difference between these indices as

$$d = \text{AFRI} - \text{WFRI}.$$

For instance, over the past year we have observed:

- **Hyperliquid:** $d \approx 14.7\%$
- **GMX V2:** $d \approx 26.2\%$

These observations suggest that lower-capacity perpetuals (e.g., on GMX V2) may yield higher returns. Figures 1 and 2 illustrate historical comparisons of AFRI and WFRI for GMX V2 and Hyperliquid, respectively.




Figure 1: Historical AFRI vs. WFRI for GMX V2. The green line represents the difference, $\Delta = \text{AFRI} - \text{WFRI}$.

2 System Architecture

The BasisOS protocol is implemented as a coordinated framework of four core smart contracts:

1. **Vault Contract:** Centralizes deposit and withdrawal management under the ERC-4626 standard [7].
2. **Strategy Contract:** Oversees portfolio rebalancing and maintains overall state management.

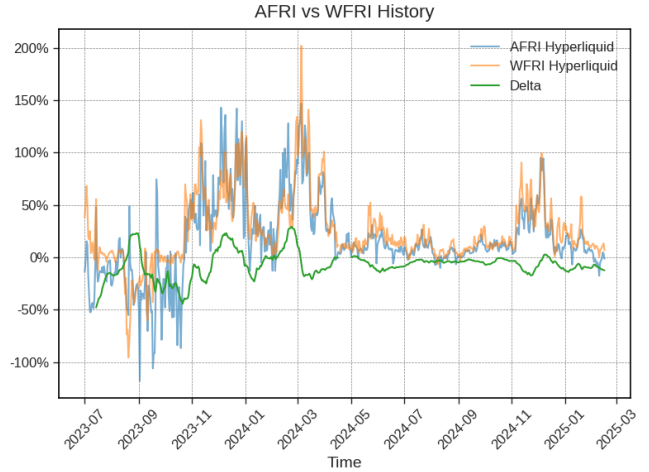


Figure 2: Historical AFRI vs. WFRI for Hyperliquid. The green line represents the difference, $\Delta = \text{AFRI} - \text{WFRI}$.

3. **Position Manager Contract:** Manages hedge positions by interfacing with Perp DEXs (e.g., Hyperliquid) using a request-response protocol.
4. **Spot Manager Contract:** Executes spot market trades to ensure alignment with hedge positions.

Off-chain keepers support each Position Manager by executing market requests, ensuring timely interaction with the corresponding Perp DEX.

3 Execution Framework

To overcome the limitations of naive execution, our managed strategy incorporates an off-chain Operator. This Operator continuously monitors market conditions and initiates transactions when spreads are favorable, thereby minimizing execution costs and slippage. Key features include:

- **Cost Efficiency:** The Operator minimizes transaction fees and slippage by guaranteeing a fixed execution spread.
- **Enhanced Spot Execution:** Integration with aggregation routers (e.g., 1Inch) improves trade execution in the spot market.

- **Batch Processing:** Simultaneous deposits and withdrawals can be batched, thereby capturing arbitrage premiums.
- **Algorithmic Timing:** The Operator algorithmically optimizes execution to take advantage of favorable market conditions.

Challenges

Despite its advantages, this approach introduces certain challenges:

- **Trustworthiness:** The off-chain Operator requires robust safeguards to mitigate trust issues.
- **Capital Utilization Delays:** Processing delays in deposits may result in temporarily suboptimal capital utilization.
- **Non-Atomic Withdrawals:** Withdrawals are processed non-atomically, necessitating additional mechanisms (e.g., Claim functions) to ensure consistency.

Despite these challenges, our managed approach significantly enhances capacity and operational efficiency.

4 Rebalancing Algorithm

The trading strategy continuously rebalances its positions to maintain the target leverage L_{target} amid market fluctuations.

Equity and Target Allocation

Let the total equity E be defined as the sum of the notional balances in the hedge and spot markets:

$$E = B_{\text{hedge}} + B_{\text{spot}},$$

where B_{hedge} and B_{spot} are the respective balances. The target allocations are then:

$$T_{\text{hedge}} = \frac{E}{1 + L_{\text{target}}}, \quad T_{\text{spot}} = \frac{E \cdot L_{\text{target}}}{1 + L_{\text{target}}}.$$

Rebalancing Triggers

Define the allocation deviations as:

$$\Delta_{\text{spot}} = T_{\text{spot}} - B_{\text{spot}}, \quad \Delta_{\text{hedge}} = T_{\text{hedge}} - B_{\text{hedge}}.$$

A rebalancing event is triggered if:

- The deviation in the hedge position exceeds a preset fraction ϵ of the corresponding spot target.
- The current leverage L_{current} falls outside the interval $[L_{\text{min}}, L_{\text{max}}]$.

5 Risk Framework

The vault's operation is governed by critical risk parameters that ensure the system's robustness. Two key metrics are defined for each asset with price trajectory $P(t)$.

Margin Treasury

The Margin Treasury (MT) is defined as the 95th percentile of the 5-minute price change distribution where price is a candlestick data per vault's asset:

$$MT = Q_{0.95}^{(5m)}$$

This threshold serves as a primary risk constraint that must not be breached under normal conditions.

Maximum Leverage

The Risk Maximum Leverage, RL_{max} , is defined such that there is only a 1% probability that an adverse price movement (without rebalancing) breaches the critical price P_{MT} .

We define P_{MT} as the price at which there remains $MT\%$ margin until the liquidation price on the exchange. In other words, if P_{liq} denotes the liquidation price, then

$$P_{MT} = P_{\text{liq}} * MT$$

Let $L(P(t))$ denote the function of the leverage at price $P(t)$, and let $Q_{0.99}^{(15m)}$ be the 99th percentile

of a 15-minute price change, given that each vault operates on a 1-minute execution cycle and the average execution latency is 2 minutes then this risk parameter takes into account extra execution delays such as Layer Zero downtimes and etc.

We then define (note that basis strategies hold short positions)

$$RL_{\max} = L \left(\frac{P_{MT}}{1 + Q_{0.99}^{(15m)}} \right).$$

For a detailed derivation of the liquidation formula for P_{liq} , refer to [3].

Assets Clustering

To calibrate risk parameters—particularly for exotic assets (e.g., memecoins) with limited historical data—statistical clustering is applied using Binance candlestick data. Key metrics include:

$$Q_{0.99}^{(5m)} = Q_{0.99} \left(\frac{P_{\text{high}}}{P_{\text{low}}} - 1 \right),$$

$$\bar{v}^{(5m)} = \frac{1}{N} \sum_{i=1}^N \frac{P_{\text{high}}^{(i)} - P_{\text{low}}^{(i)}}{P_{\text{open}}^{(i)}},$$

$$v_{0.99}^{(5m)} = Q_{0.99} \left(\frac{P_{\text{high}} - P_{\text{low}}}{P_{\text{open}}} \right),$$

with analogous metrics defined for 15-minute candles and the maximum leverage reported by the Hyperliquid exchange. A k -Nearest Neighbors clustering algorithm is used to group assets based on these features. Figures 3 and 4 illustrate the clustering results.

Results

Using historical data and the Hyperliquid liquidation price formula implemented in Fractal[2], we estimated RL_{\max} for selected tokens. The liquidation price is computed as:

$$P_{liq} = P - \frac{\text{side} \cdot M_{\text{avail}}}{\text{PositionSize} \left(1 - \frac{\text{side}}{L_{\text{maint}}} \right)}.$$

Table 1 summarizes the RL_{\max} values.

As mentioned earlier, there is insufficient historical data for young assets to reliably estimate




Figure 3: $Q_{0.99}^{(5m)}$ versus $\bar{v}^{(5m)}$ across asset clusters.

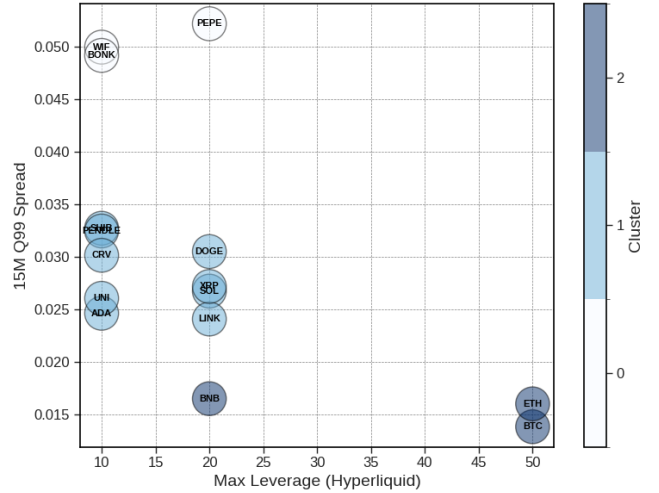


Figure 4: Maximum leverage versus $Q_{0.99}^{(15m)}$ across asset clusters.

Ticker	RL_{\max}
BTC	13.1
ETH	12.7
SOL	9.4
DOGE	9.0
PEPE	7.5

Table 1: RL_{\max} values for selected tokens derived from historical data.

their risk parameters. Therefore, we will use the risk parameters determined for PEPE as a conservative benchmark for all tokens lacking sufficient data to optimize custom risk parameters.

6 Strategy Capacity

Vault capacity is defined by constraints on both the spot and hedge sides, ensuring that positions can be closed with acceptable slippage while maintaining sufficient market depth.

Spot Capacity

Spot capacity is limited by:

1. **Liquidity Constraint:**

$$CAP_{\text{liquidity}} = T_l \times \$\text{AvgPoolTVL},$$

where $\$\text{AvgPoolTVL}$ is the average total value locked for the asset and T_l (e.g., 5%) is the maximum fraction allowed.

2. **Slippage Constraint:**

$$\text{slpg}(CAP_{\text{slippage}}) < T_s,$$

where T_s is the maximum permissible slippage.

Thus, the effective spot capacity is:

$$CAP_{\text{spot}} = \min(CAP_{\text{liquidity}}, CAP_{\text{slippage}}).$$

Hedge Capacity

Hedge capacity is determined by:

$$CAP_{\text{hedge}} = T_{oi} \times \$\text{AvgOI},$$

where $\$\text{AvgOI}$ is the average open interest and T_{oi} (e.g., 10%) is the maximum allowed fraction [5].

Effective Vault Capacity

The overall vault capacity is given by:

$$\text{Vault Capacity} = \min(CAP_{\text{spot}}, CAP_{\text{hedge}}).$$

These constraints are critical to managing liquidity risk in basis trading [6].

7 Simulation

Fractal Framework

The Fractal Framework is a modular research library designed for developing and backtesting DeFi strategies. It models financial entities (e.g., spot trading, liquidity pools, and hedging) as state functions, facilitating rigorous simulation and analysis. Further details are available in the Fractal documentation [2].

Methodology

Our simulation methodology integrates:

1. **Historical-Based Synthetic Data:** A sliding window technique is used to extract multiple sub-trajectories from historical data.
2. **Monte Carlo Simulations:** Synthetic market scenarios are generated based on historical statistics to introduce stochasticity and test robustness.

Objective Function

Let $\{APY_i\}_{i=1}^N$ denote the annualized yields over N simulation trajectories, with the mean APY given by

$$\bar{A} = \frac{1}{N} \sum_{i=1}^N APY_i.$$

Let $\{DD_i\}_{i=1}^N$ denote the maximum drawdowns observed, and define the 5% quantile drawdown DD^{q5} as

$$DD^{q5} = \inf \left\{ x \in \mathbb{R} \left| \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{DD_i \leq x\} \geq 0.05 \right. \right\}.$$

We penalize drawdown risk with the weight α via

$$D_1 = 1 - \alpha DD^{q5}, \quad J_1 = \frac{\bar{A}}{D_1}.$$

For leverage asymmetry, define

$$\Delta_{\text{max}} = L_{\text{max}} - L_{\text{target}}, \quad \Delta_{\text{min}} = L_{\text{target}} - L_{\text{min}},$$

$$\Delta = |\Delta_{\max} - \Delta_{\min}|,$$

and with the weight β let

$$D_2 = 1 - \beta \Delta.$$

The overall objective function is then given by

$$F = \frac{\bar{A}}{(1 - \alpha DD^{q5})(1 - \beta \Delta)}.$$

8 Results

Simulations were conducted using Binance historical data, covering diverse funding rate regimes (both positive and negative) and various market conditions. The time range was chosen based on the maximum available history. Note that BTC and ETH used 1-hour candlesticks, while DOGE and PEPE employed 5-minute candlesticks due to high volatility. In the following subsections, we present results for BTC, ETH, DOGE, and PEPE.

The average execution cost per rebalancing was chosen empirically by evaluating the historical spreads for these assets and considering the current fees on platforms such as GMX V2, Hyperliquid, Uniswap V2, Uniswap V3, and 1inch.

Asset	Execution Cost (%)
BTC	0.2
ETH	0.2
DOGE	0.4
PEPE	0.5

Table 2: Execution cost per rebalancing for different assets.

BTC Results

Figures 5–9 display the BTC objective function surface, the simulated maximum drawdown distribution, and the backtesting performance. Table 3 lists the top parameter sets.

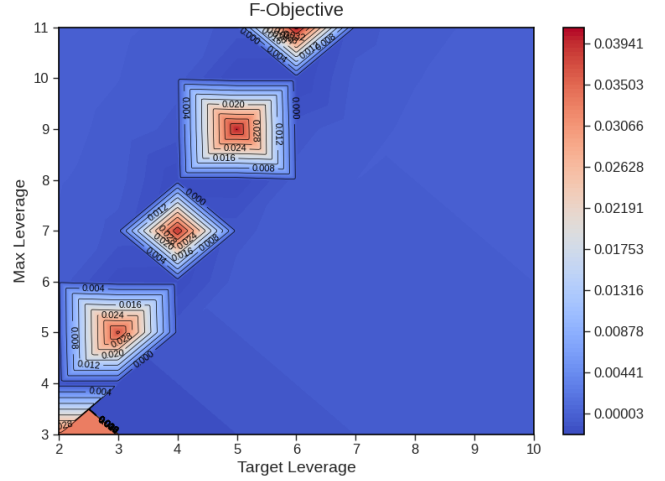


Figure 5: Objective function $F_{\text{BTC}}(L_{\text{target}}, L_{\text{max}})$ for BTC.

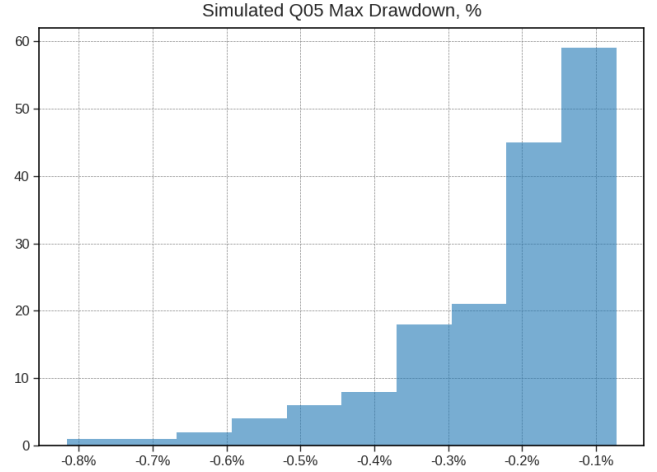


Figure 6: Simulated 5% quantile of BTC maximum drawdown, D_{BTC}^{q5} .

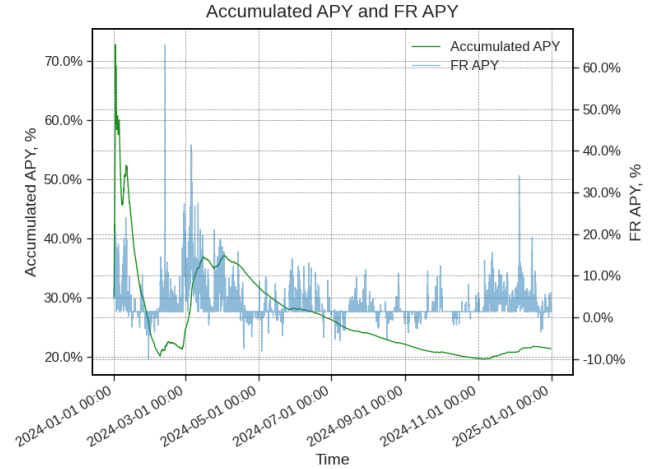


Figure 7: BTC: Accumulated APY over time.

Table 3: Top 5 BTC Parameter Sets

L_{\min}	L_{target}	L_{\max}	Objective	Backtest Acc. PNL	Backtest APY	Backtest MAX DD	Simulated AVG APY
1.0	6.0	11.0	0.0417	13.87%	4.62%	-0.25%	4.33%
1.0	5.0	9.0	0.0405	13.55%	4.52%	-0.20%	4.22%
1.0	4.0	7.0	0.0392	13.04%	4.34%	-0.16%	4.10%
2.0	6.0	10.0	0.0383	13.08%	4.36%	-0.24%	4.13%
1.0	3.0	5.0	0.0370	12.52%	4.17%	-0.14%	3.89%

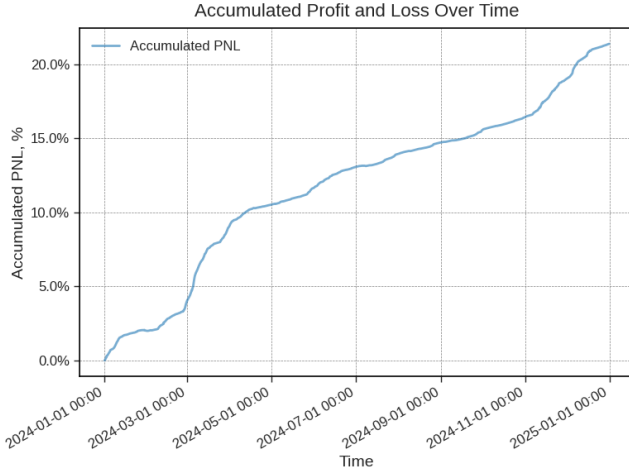


Figure 8: BTC: Accumulated Profit and Loss over time.

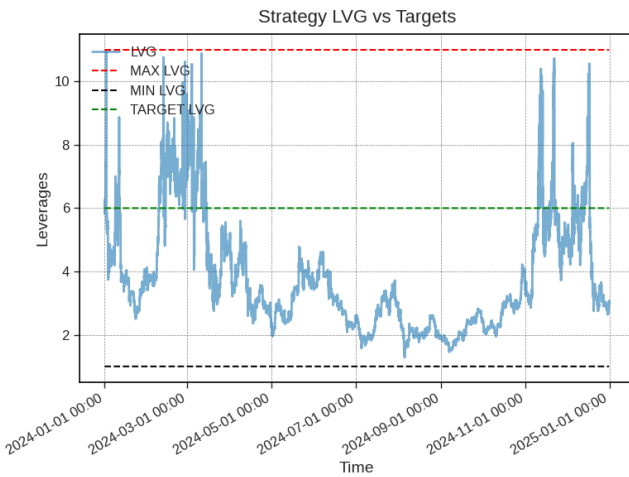


Figure 9: BTC: Strategy leverage compared to target and bounds.

ETH Results

Figures 10–14 present the ETH objective function, the simulated drawdown distribution, and the backtesting performance. Table 4 shows the top parameter sets.




Figure 10: Objective function $F_{\text{ETH}}(L_{\text{target}}, L_{\max})$ for ETH.

DOGE Results

For DOGE, an asset with higher volatility but substantial liquidity, Figures 15 and 16 show the objective function surface and the simulated 5% drawdown distribution. Table 5 lists the top parameter sets.

PEPE Results

PEPE, an exotic memecoin with substantially higher volatility, is evaluated next. Figures 20 and 21 display the objective function surface and the simulated 5% drawdown distribution. Table 6 provides the top parameter sets.

Table 4: Top 5 ETH Parameter Sets

L_{\min}	L_{target}	L_{\max}	Objective	Backtest Acc. PNL	Backtest APY	Backtest MAX DD	Simulated AVG APY
1.0	5.0	9.0	0.0369	12.06%	4.02%	-1.05%	4.07%
1.0	4.0	7.0	0.0354	11.64%	3.88%	-0.95%	3.92%
1.0	3.0	5.0	0.0332	10.61%	3.53%	-1.18%	3.66%
2.0	6.0	10.0	0.0321	11.48%	3.82%	-1.71%	3.76%
2.0	5.0	8.0	0.0319	11.16%	3.72%	-1.64%	3.68%

Table 5: Top 5 DOGE Parameter Sets

MIN LVG	TRGT LVG	MAX LVG	Objective	Backtest Acc. PNL	Backtest APY	Backtest MAX DD	Simulated AVG APY
1.0	4.0	7.0	0.0499	4.47%	4.46%	-0.74%	5.48%
1.0	3.0	5.0	0.0383	4.80%	4.79%	-0.61%	4.54%
1.0	2.0	3.0	0.0347	3.39%	3.38%	-0.52%	3.98%
2.0	4.0	6.0	0.0148	1.28%	1.28%	-1.82%	1.79%
4.0	5.0	7.0	0.0039	-10.36%	-10.34%	-11.05%	-11.26%

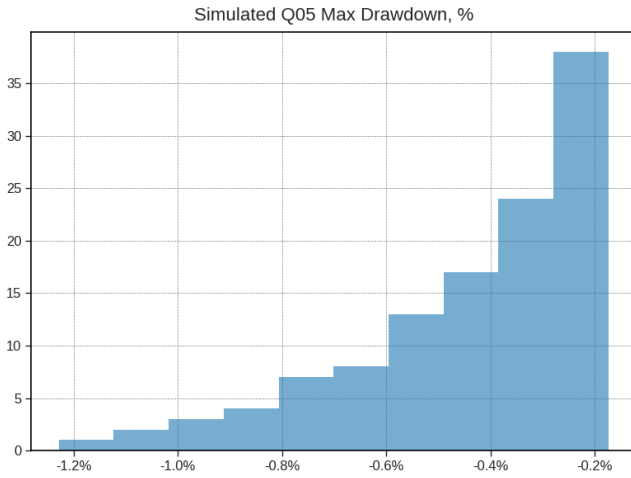


Figure 11: Simulated 5% quantile of ETH maximum drawdown, D_{ETH}^{q5} .

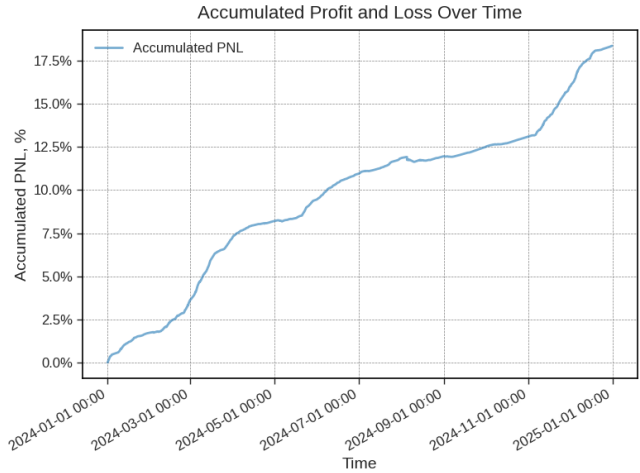


Figure 13: ETH: Accumulated Profit and Loss over time.




Figure 12: ETH: Accumulated APY over time.




Figure 14: ETH: Strategy leverage compared to target and bounds.

Table 6: Top 5 PEPE Parameter Sets

MIN LVG	TRGT LVG	MAX LVG	Objective	Backtest Acc. PNL	Backtest APY	Backtest MAX DD	Simulated AVG APY
1.0	3.0	5.0	0.0438	6.21%	6.19%	-0.90%	5.16%
1.0	2.0	3.0	0.0308	3.05%	3.04%	-1.48%	3.61%
3.0	5.0	6.0	0.0051	-13.60%	-13.57%	-13.89%	-15.02%
3.0	4.0	6.0	0.0047	-12.86%	-12.84%	-13.14%	-13.97%
2.0	4.0	5.0	0.0017	-4.82%	-4.81%	-6.93%	-4.21%

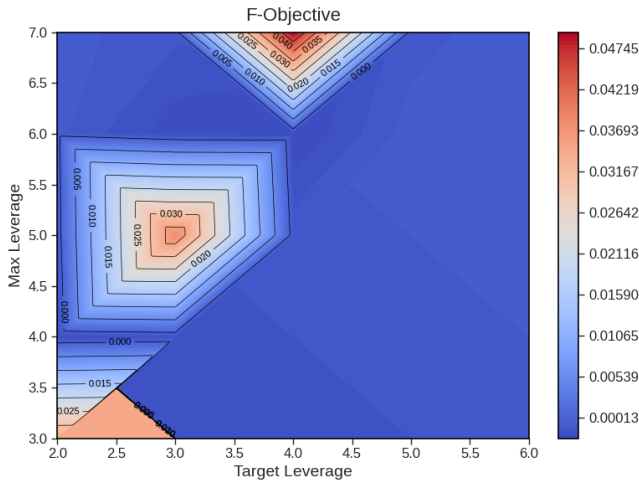


Figure 15: DOGE: Objective function $F_{\text{DOGE}}(L_{\text{target}}, L_{\text{max}})$.

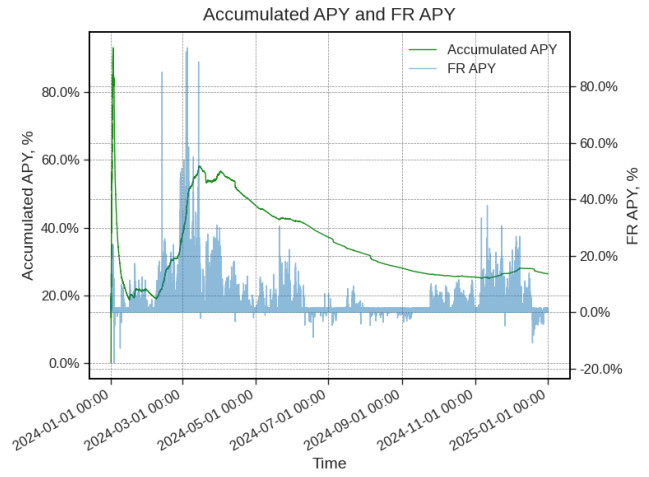


Figure 17: DOGE: Accumulated APY (green line) and Funding Rate APY (blue bars).

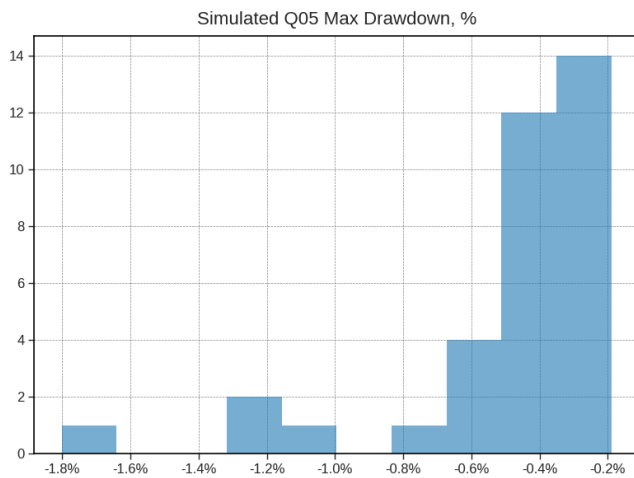


Figure 16: DOGE: Simulated 5% maximum drawdown distribution.

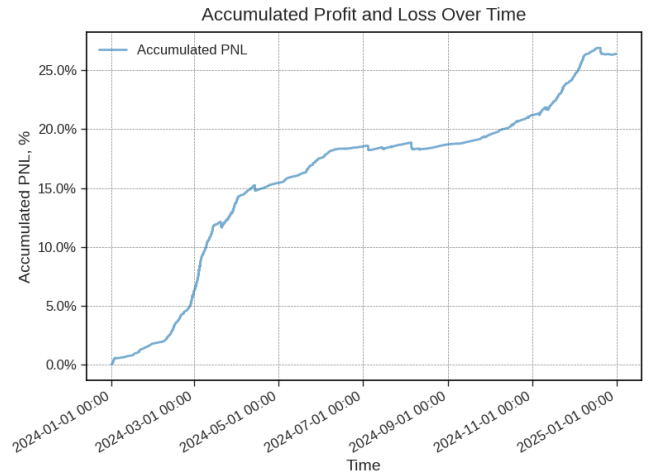


Figure 18: DOGE: Accumulated Profit and Loss over time.




Figure 19: DOGE: Strategy leverage compared to target and bounds.

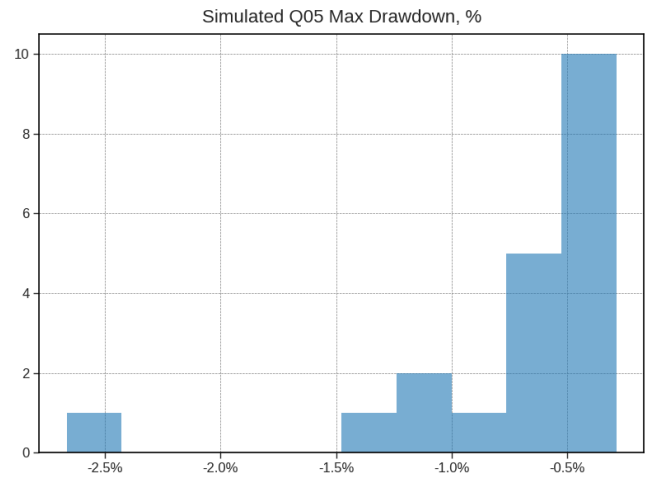


Figure 21: PEPE: Simulated 5% maximum draw-down distribution.

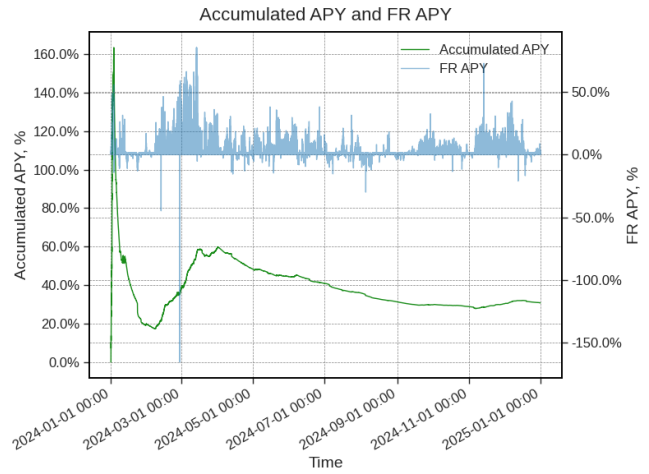


Figure 22: PEPE: Accumulated APY (green line) and Funding Rate APY (blue bars).




Figure 20: PEPE: Objective function $F_{PEPE}(L_{target}, L_{max})$.

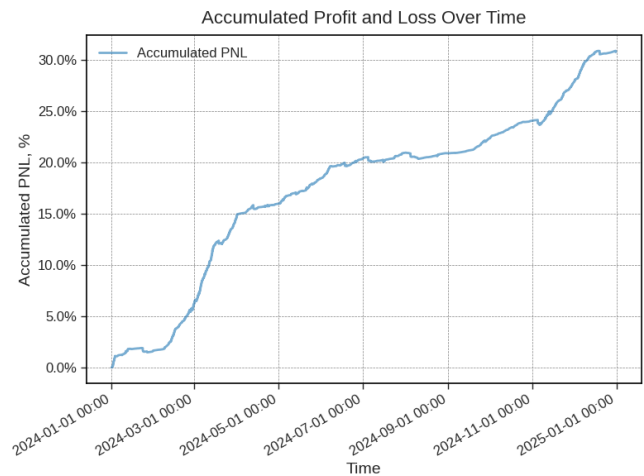


Figure 23: PEPE: Accumulated Profit and Loss over time.

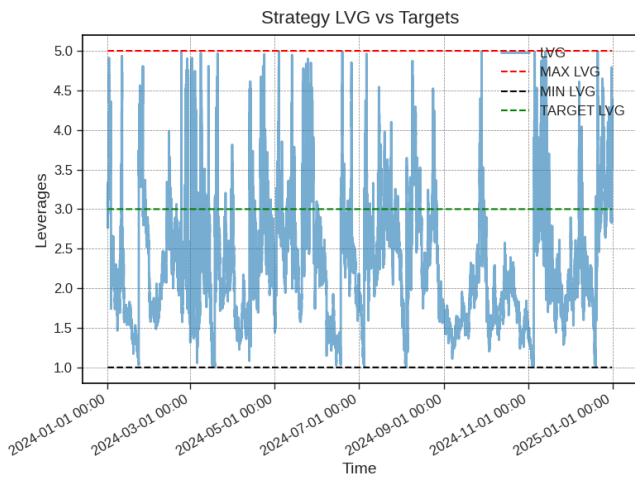


Figure 24: PEPE: Strategy leverage compared to target and bounds.

Simulation and Backtest Summary

In summary, simulation results indicate that more volatile assets yield less optimal parameter solutions due to a restricted leverage range. For exotic assets, it is crucial to balance the frequency of rebalancing with maintaining high vault utilization to achieve high APY. The following table summarizes the final Hyperliquid backtest results for the 2024–2025 period using the best parameter sets:

Asset	APY (%)	Max Drawdown (%)
BTC	21.38	-0.063
ETH	18.37	-0.256
DOGE	27.56	-0.396
PEPE	30.82	-0.507

9 Benchmarks

To evaluate the performance of the BasisOS vault strategies against existing USD-based DeFi solutions, we conducted comparative backtests over the same 2024–2025 time horizon. Figures 25 and 26 illustrate the annualized percentage yield (APY) profiles of:

- **BasisOS Optimized Parameters on Hyperliquid:** BTC, ETH, DOGE, and PEPE.

- **USD-Based DeFi Strategies:** AAVE USDC APY, Ethena APY, USDX APY, Resolv stUSDr APY, and BasisOS ETH APY.

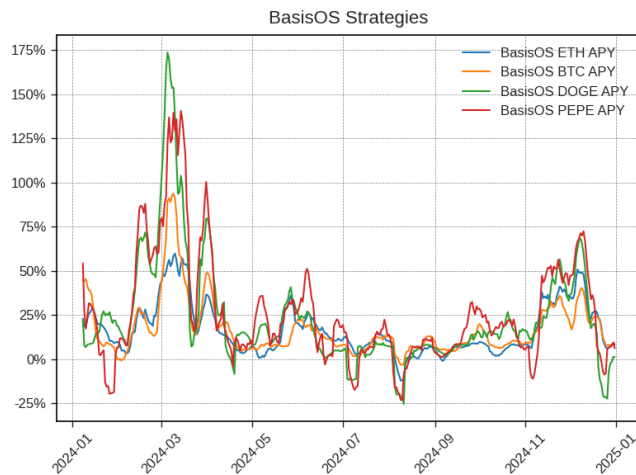


Figure 25: BasisOS strategies (BTC, ETH, DOGE, PEPE) backtested on Hyperliquid over the 2024–2025 period.

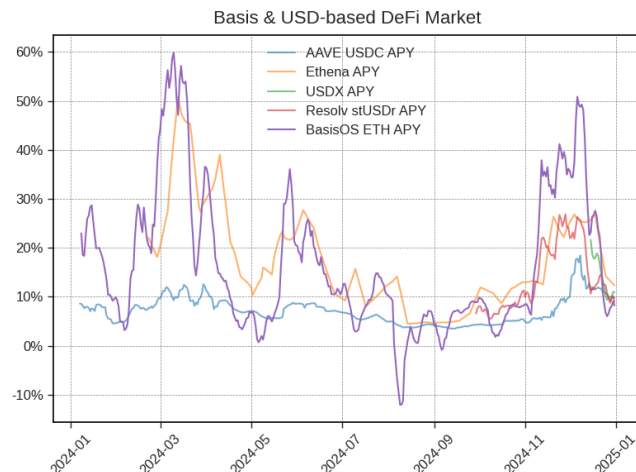


Figure 26: Comparison of BasisOS ETH APY versus leading USD-based DeFi strategies (AAVE USDC, Ethena, USDX, and Resolv stUSDr) over the 2024–2025 period.

Despite increased volatility, assets like DOGE and PEPE with optimized capital utilization can generate superior returns by capturing elevated funding rates. Both blue-chip assets (such as BTC and ETH) and memecoins (such as DOGE and PEPE) that were backtested outperform all existing benchmarks.

References

- [1] Ethena documentation. <https://docs.ethena.fi/>. Accessed: 2025-02-17.
- [2] Fractal framework. <https://logarithm-labs.gitbook.io/fractal>. Accessed: 2025-02-17.
- [3] Hyperliquid liquidation formula. <https://docs.hyperliquid.org/>. Accessed: 2025-02-17.
- [4] Perphub. <https://perphub.com/>. Accessed: 2025-02-17.
- [5] Logarithm Finance. Logarithm basis funding risks. <https://medium.com/@logarithm.finance/logarithm-basis-funding-risks-288474b19277>. Accessed: 2025-02-17.
- [6] Logarithm Finance. Logarithm basis liquidity risks. <https://medium.com/@logarithm.finance/logarithm-basis-liquidity-risks-c9e6ea2ecbfe>. Accessed: 2025-02-17.
- [7] Joey Santoro, t11s, Jet Jadeja, Alberto Cuesta Cañada, and Señor Doggo. Erc-4626: Tokenized vaults. Technical report, Ethereum Improvement Proposals, Dec 2021. Created 2021-12-22. Accessed: February 17, 2025. <https://eips.ethereum.org/EIPS/eip-4626>.