

7514718-online-learning-of-arbitrage-capacity-thresholds

SUMMARY: This paper investigates the threshold at which the funding-rate mean-reversion mechanism for crypto perpetual funding rates stops working. It uses hourly Binance USDT-M futures data for BTC, ETH, and SOL from January 2021 to December 2024. The study identifies the arbitrage capacity threshold, where a drop in USDe supply from \$15 billion to \$4.3 billion reduces the strategy Sharpe ratio from 2.805 to 0.319. The paper validates online logistic regression with inference-time adaptation as an effective filter for non-stationary markets, outperforming random filtering in 8 out of 9 parameter combinations. Cross-validation holds for BTC and ETH, but not for SOL. The study also examines the impact of USDe supply on strategy performance, showing a significant difference in Sharpe ratios when USDe supply exceeds or falls below \$3 billion. The findings suggest that online logistic regression can improve the Sharpe ratio for BTC and ETH, but not for SOL, highlighting the importance of underlying signal quality.

METHOD: The study uses hourly perpetual futures data for BTC, ETH, and SOL from Binance USDT-M futures, covering January 2021 to December 2024. The primary signal is a funding-rate mean-reversion event, defined by a Z-score threshold of -1.85 . The events are labeled using a Triple-Barrier scheme with a 24-hour holding period, and take-profit and stop-loss thresholds are set at 2.0 ATR and 1.0 ATR, respectively. The feature set includes seven regime variables: 24-hour and 168-hour realized volatility, their ratio, 24-hour and 168-hour momentum, 24-hour range, and 30-day absolute funding rate median. The online logistic regression filter is updated at every event, with parameters adjusted using an adaptive learning rate. The study also examines the impact of USDe supply on strategy performance, using the Ethena USDe supply series as a proxy for arbitrage capital capacity.

KEY FINDINGS:

- The funding-rate mean-reversion strategy for BTC, ETH, and SOL has a Sharpe ratio of 1.248, -0.084 , and -0.879 , respectively.
- When USDe supply drops from \$15 billion to \$4.3 billion, the strategy Sharpe ratio drops from 2.805 to 0.319.
- Online logistic regression with inference-time adaptation improves the Sharpe ratio for BTC and ETH, reducing drawdown from -14.52% to -4.72% .
- The gain in Sharpe ratio comes from genuine signal filtering, as demonstrated by ablation tests against 50 random-filter simulations.
- The study finds that the gain in Sharpe ratio does not hold for SOL, suggesting that the effectiveness of machine learning filters depends on the quality of underlying signals.

IMPLICATION FOR TRADING: The findings suggest that online logistic regression can be an effective tool for filtering non-stationary markets, particularly for BTC and ETH. However, the effectiveness may vary depending on the underlying signal quality, as demonstrated by the poor performance of the strategy for SOL. Practitioners can use this information to improve their trading strategies by incorporating machine learning filters, such as online logistic regression, to identify and exploit mean-reversion opportunities in crypto markets. Additionally, the study highlights the importance of understanding the impact of USDe supply on strategy performance, which can inform trading decisions and risk management.

Online Learning of Arbitrage Capacity Thresholds

Evidence from Crypto Perpetual Funding Rates

September 23, 2026

Abstract

Funding rate stability depends on arbitrage capital capacity. A funding-rate mean-reversion strategy delivers Sharpe ratios of 1.248 (BTC), -0.084 (ETH), and -0.879 (SOL). Online logistic regression with inference-time adaptation raises BTC’s Sharpe to 1.570 and cuts drawdown from -14.52% to -4.72% . Ablation against 50 random-filter simulations and across 9 parameter combinations confirms the gain comes from genuine signal filtering. Cross-asset validation holds for BTC and ETH (ETH improves from -0.084 to 1.006) but not for SOL.

1 Introduction

On 10 October 2025, Hyperliquid absorbed \$641 million in forced selling within a single hour. The backstop mechanism absorbed \$576 million (62.6%), and only \$64 million reached the public order book. This was not a black swan. It was the predictable consequence of arbitrage capital contraction.

This paper asks a single question: when does the funding-rate mean-reversion mechanism stop working? The answer is not “when market maturity is low,” as the existing literature suggests [1]. The answer is “when arbitrage capital capacity falls below a threshold.”

I make three contributions. First, I identify the arbitrage capacity threshold in crypto perpetual funding rates. When USDe supply falls from \$15 billion to \$4.3 billion, strategy Sharpe drops from 2.805 to 0.319. Second, I validate online logistic regression with inference-time adaptation as an effective filter for non-stationary markets. Across 9 parameter combinations, it outperforms random filtering in 8 out of 9. Third, I provide cross-asset heterogeneity evidence: the gain holds for BTC and ETH but not for SOL, suggesting ML filtering effectiveness depends on underlying signal quality.

This paper builds on prior work on strategic response direction [6] and on the disappearance of hedging relationships [7]. The first establishes that the direction of strategic response determines the direction of threshold movement. The second documents that broken relationships do not automatically restore. The USDe supply series is used as a proxy for arbitrage capital capacity, following the shock-identification approach in [3].

The rest of the paper proceeds as follows. Section 2 describes data and method. Section 3 reports cross-sectional and time-series results. Section 4 presents the ML ablation and robustness

tests. Section 5 discusses mechanisms and boundaries. Section 6 concludes.

2 Data and Method

2.1 Data

I use hourly perpetual futures data for BTC, ETH, and SOL from Binance USDT-M futures, covering January 2021 to December 2024. The sample contains 35,064 hourly observations per coin. Funding rates are 8-hour settlements, aligned to hourly bars via forward fill. I also use the Ethena USDe supply series (34 data points from December 2023 to August 2026) as a proxy for arbitrage capital capacity.

2.2 Signal construction

The primary signal is a funding-rate mean-reversion event. Define the funding rate Z-score as

$$z_t = \frac{f_t - \bar{f}_{720}}{\hat{\sigma}_{720}}, \quad (1)$$

where f_t is the funding rate at hour t , and $\bar{f}_{720}$ and $\hat{\sigma}_{720}$ are the 720-hour rolling mean and standard deviation. A jump-down event occurs when $z_t < -1.85$ and $z_{t-1} > -0.5$.

2.3 Labeling and execution

Each event is labeled using a Triple-Barrier scheme with a 24-hour holding period. Take-profit and stop-loss thresholds are set at 2.0 ATR and 1.0 ATR respectively. Round-trip transaction cost is 0.06%.

2.4 Online logistic regression

The ML filter is an online logistic regression with inference-time adaptation. Parameters are updated at every event, not frozen:

$$p_t = \sigma(x_t^\top w_t + b_t), \quad (2)$$

$$w_{t+1} = w_t - \eta_t [(p_t - y_t)x_t + \lambda w_t], \quad (3)$$

$$b_{t+1} = b_t - \eta_t (p_t - y_t), \quad (4)$$

$$\eta_{t+1} = \eta_t \cdot \gamma. \quad (5)$$

Baseline parameters are $\eta_0 = 0.05$, $\lambda = 0.001$, $\gamma = 0.999$. The feature set contains seven regime variables: 24-hour and 168-hour realized volatility, their ratio, 24-hour and 168-hour momentum, 24-hour range, and 30-day absolute funding rate median.

3 Empirical Results

3.1 Cross-sectional differences

Table 1 reports the baseline strategy performance. Sharpe ratios vary substantially across coins: BTC at 1.248, ETH at -0.084 , SOL at -0.879 . The out-of-sample degradation for BTC is -15.5% , confirming the result is not driven by parameter mining.

Table 1: Baseline strategy performance, $z_{\text{down}} = -1.85$, $\text{hold} = 24$ hours.

Coin	Events	Annual return	Sharpe	Max drawdown
BTC	75	19.33%	1.248	-14.52%
ETH	57	-1.94%	-0.084	-24.71%
SOL	86	-13.15%	-0.879	-54.51%




Figure 1: Cross-sectional Sharpe ratios: baseline versus online logistic regression. The online filter raises BTC and ETH but fails to rescue SOL.

3.2 Time-series regime shift

Table 2 reports strategy performance conditional on USDe supply. When USDe exceeds \$3 billion, Sharpe is 2.805. When it falls below \$3 billion, Sharpe is 0.319. The three-way split shows the same asymmetry: high-USDe on up-days delivers 4.703, while high-USDe on down-days delivers -1.53 . The regime-stratified reporting used here follows the three-pillar framework recommended in [?].

Table 2: Conditional performance by USDe supply regime.

Regime	Events	Annual return	Sharpe
High USDe ($> \$3\text{B}$)	9	56.94%	2.805
Low USDe ($\leq \$3\text{B}$)	6	1.05%	0.319
High USDe, up-days	6	86.61%	4.703
High USDe, down-days	7	-15.55%	-1.53




Figure 2: Sharpe ratio by USDe supply regime and daily direction. The asymmetry between up-days and down-days is the core mechanism.

4 Machine Learning Validation

4.1 Online logistic regression results

Table 3 reports the ML filtering results. Online logistic regression at threshold 0.35 delivers Sharpe 1.570 for BTC, up from baseline 1.251, with drawdown reduced from -14.52% to -4.72% . Seven alternative ML methods (GradientBoosting, HMM, volatility prediction, ProSP, BOCPD, and Bayesian online LR) deliver no stable gain.

Table 3: ML filtering comparison, BTC.

Method	Sharpe
Baseline	1.251
Online LR, th= 0.35	1.570
Online LR, th= 0.40	1.558
GradientBoosting	1.196
HMM regime routing	out-of-sample failure
Bayesian online LR	1.528

4.2 Ablation against random filtering

To verify that the gain comes from genuine filtering rather than from simply reducing trade count, I run 50 random-filter simulations retaining the same proportion of events (69.3%). The mean Sharpe is 1.023 and the 95th percentile is 1.378. The online logistic regression Sharpe of 1.570 exceeds the 95th percentile.

4.3 Parameter robustness

Table 4 reports the 3×3 sweep over z_{down} and holding period. Eight of nine combinations pass the random-filter test. The only failure occurs at $z = -2.0$ and $\text{hold} = 36$.




Figure 3: Random-filter distribution (50 trials) versus online logistic regression. The vertical line marks the online LR Sharpe of 1.570.

Table 4: Joint sweep: $z_{\text{down}} \times \text{holding period}$. Each cell reports online Sharpe and pass/fail against the random-filter 95th percentile.

	hold = 12	hold = 24	hold = 36
$z = -1.65$	1.229 (pass)	1.440 (pass)	1.295 (pass)
$z = -1.85$	1.012 (pass)	1.570 (pass)	1.546 (pass)
$z = -2.00$	1.112 (pass)	1.190 (pass)	0.984 (fail)

4.4 Cross-asset validation

The online filter works for BTC and ETH but not for SOL. For ETH, Sharpe rises from -0.084 to 1.006 ; for SOL, Sharpe remains negative at -0.515 and fails the random-filter test. A probability-sorted portfolio (ProSP) experiment, following [?], produced zero cross-sectional spread (top and bottom signals both at 0.302 win rate).

5 Mechanism and Boundaries

5.1 Arbitrage capacity as the missing variable

The cross-sectional and time-series results point to a single mechanism: funding rate stability depends on arbitrage capital capacity. When capacity is sufficient (USDe above \$3 billion), the mean-reversion mechanism is anchored and the strategy is profitable. When capacity contracts (USDe below \$3 billion), the anchor breaks.

This interpretation is consistent with three strands of literature. Haddad, Huebner, and Loualiche [2] show that strategic reaction offsets only two-thirds of shocks in U.S. equity markets. BitMEX Research [1] documents that funding rates cluster at 0.01% for 92% of the time, maintained by arbitrage capital. Ethena’s USDe supply contraction from \$15 billion to \$4.3 billion provides a direct measure of capacity withdrawal in the crypto perpetual market.

The capacity-threshold mechanism developed here extends the framework in [6] from equity markets to crypto perpetual funding rates. It also connects to the finding in [7] that hedging relationships can break without recovery.

5.2 Online learning as a proxy for capacity presence

A second implication is methodological. The effectiveness of online logistic regression is itself a proxy for whether arbitrage capital is present. In BTC and ETH, the filter succeeds (Sharpe above the random-filter 95th percentile), implying that the signal contains identifiable structure. In SOL, the filter fails, implying the signal is noise. This provides an independent measure of capacity presence, not requiring direct measurement of arbitrage concentration.

5.3 Boundaries

Three boundaries should be noted. First, the sample is small (75 BTC events, 57 ETH events, 86 SOL events), and the confidence intervals around the Sharpe estimates are wide. Second, the empirical strategy is descriptive rather than causal; the 2025-10 Ethena collapse offers a potential quasi-natural experiment but is outside the current sample window. Third, the ML filter is the simplest form of inference-time adaptation; comparison against frontier methods such as PULSE [4] and DREAM [5] is left for future work.

6 Conclusion

This paper identifies an arbitrage capacity threshold in crypto perpetual funding rates. Using a funding-rate mean-reversion strategy and an online logistic regression filter with inference-time adaptation, I show that the mechanism is effective for BTC and ETH but fails for SOL. The gain is verified by ablation against random filtering and remains robust across 9 parameter combinations. The effectiveness of the online filter itself serves as a proxy for arbitrage capital presence.

References

- [1] BitMEX Research. 2026. “Three Sources of Funding-Rate Alpha.” Q2 2026 Derivatives Report.
- [2] Haddad, V., P. Huebner, and E. Loualiche. 2025. “How Competitive Is the Stock Market? Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing.” *American Economic Review* 115 (3): 975–1018.
- [3] International Monetary Fund. 2026. “Stablecoin Shocks and Macro-Financial Spillovers.” IMF Working Paper 2026/044.
- [4] Cartea, Á., G. Duran-Martin, and L. Sánchez-Betancourt. 2026. “Detecting Toxic Flow.” *Quantitative Finance* 26 (4): 541–561.
- [5] Karl, B. 2026. “DREAM: Dynamic Recall and Elastic Adaptive Memory on Machine Learning.” Zenodo preprint.
- [6] Shek, H. 2026. “Direction of Feedback: Passive Investing and the Stability Threshold.” SSRN Working Paper. DOI: 10.2139/ssrn.7458346.

- [7] Shek, H. 2026. “The Vanishing Hedge: The Term Premium and Financial Conditions Transmission after 2010.” SSRN Working Paper. DOI: 10.2139/ssrn.7458558.