

Limits to coordination: Funding drag in crypto perpetual funding rates

Helvin Shek
hs3710@nyu.edu

September 27, 2026

Abstract

Extreme negative funding rates in cryptocurrency perpetual futures are widely interpreted as reversal signals, but the speed of the subsequent correction varies by an order of magnitude. This paper introduces **limits to coordination**, a new form of arbitrage constraint operating alongside the capital constraints of Shleifer and Vishny (1997) and the funding-liquidity constraints of Brunnermeier and Pedersen (2009). I document **funding drag**: the speed of parity restoration after a funding-rate extreme is dominated by realized volatility. Using 1,696 events across 18 Binance perpetual contracts, high-volatility events take a median of 27.0 hours to correct, against 3.0 hours for low-volatility events (log-rank $\chi^2 = 147.25$, $p < 0.0001$). A one-standard-deviation increase in realized volatility lowers the hazard of correction by approximately 41% (HR = 0.5935; naive SE = 0.0549; contract-clustered SE = 0.0675; year-month-clustered SE = 0.0784). A placebo test on volatility-matched non-event hours shows that only extreme funding-rate events are slow to correct (27.0h vs 10.0h, $\chi^2 = 88.04$, $p < 10^{-20}$). Four accelerated-failure-time specifications yield time ratios of 2.32–3.72. A recursive-partitioning tree places the strongest heterogeneity at the volatility ratio vol24/vol168, with correction time spanning a 10.6-fold range from 11.53 to 122.62 hours. The effect is present both before and after the FTX collapse (HR 0.570 vs 0.702). The evidence is consistent with a coordination-failure interpretation, described here as a robust empirical regularity rather than as a structural claim about the mechanism.¹

Keywords: perpetual futures, funding rates, limits to coordination, funding drag, survival analysis, realized volatility

JEL Classification: G12, G13, G14, G15, G23, C58

1 Introduction

When the funding rate on a cryptocurrency perpetual futures contract falls to an extreme negative value, short sellers are paying long holders to maintain their positions. A trader who buys the perpetual and sells spot receives the funding payment while holding a delta-neutral position.

¹The main text and main tables use the current 1,696-event replication. The supplementary tables (S1–S7) report the original 1,699-event estimates from the paper’s first version; the two samples differ only in the listing dates used for APTUSDT, ARBUSDT, and OPUSDT. The three-event difference does not affect any point estimate by more than 0.1 standard errors nor any qualitative conclusion.

Standard arbitrage logic implies that this trade should push the perpetual price back toward its pre-event level.

But the reversal does not always occur, and when it does, it does not always occur at the same speed. The same funding-rate extreme can correct in three hours or in three days. The literature on limits to arbitrage is large, but it has not been applied to funding-rate correction speed. The present paper fills this lacuna.

This paper studies the speed of correction after funding-rate extreme events. I use 1,696 events across 18 perpetual contracts on Binance and a Kaplan-Meier survival framework to estimate how long prices take to return to within a fixed percentage of their pre-event level. The sample spans four distinct market regimes: the 2021 bull market and correction, the 2022 bear market, the 2023 recovery, and the 2024 bull market.

I document a new empirical regularity that I call *funding drag*, and introduce *limits to coordination* as its underlying mechanism. High-volatility events take a median of 27.0 hours to correct; low-volatility events take 3.0 hours.

The paper makes four contributions. First, it introduces *limits to coordination*, a new form of arbitrage constraint operating alongside the capital constraints of Shleifer and Vishny (1997) and the funding-liquidity constraints of Brunnermeier and Pedersen (2009). Second, it documents *funding drag*, the empirical pattern through which this constraint manifests. Third, a placebo test on volatility-matched non-event hours shows that only extreme funding-rate events are slow to correct. Fourth, four accelerated-failure-time specifications yield time ratios of 2.32–3.72, all above one.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data and empirical design. Section 4 presents the main results. Section 5 discusses mechanisms and implications. Section 6 concludes.

2 Related Literature

2.1 Limits to arbitrage

Shleifer and Vishny (1997) show that when arbitrageurs face capital constraints, price deviations from fundamental value can persist. Brunnermeier and Pedersen (2009) link market liquidity to funding liquidity. Makarov and Schoar (2020) document large and persistent arbitrage spreads across crypto exchanges. The evidence here suggests that realized volatility is also a determinant of arbitrage effectiveness.

2.2 Coordination failure

Diamond and Dybvig (1983) show that demand deposit contracts have multiple equilibria, one of which is a bank run. Morris and Shin (1998) show that noisy signals collapse the multiplicity to a unique equilibrium. The evidence here is consistent with this framework, but the mechanism is not directly identified.

2.3 Information acquisition and model uncertainty

Grossman and Stiglitz (1980) show that if prices already reflect all information, no one has an incentive to collect it. Epstein and Schneider (2003) axiomatize recursive multiple-priors; Hansen and Sargent (2001) parameterize the same idea using relative entropy. Wager and Athey (2018) provide the estimation framework for the heterogeneity analysis.

2.4 Crypto market microstructure

Willeboordse (2026) shows that stablecoin-linked triangles on Kraken are persistently wider than pure-fiat triangles. Pindza (2026) shows that leverage and volatility regime dominate liquidation hazard in CEX-DEX funding-rate arbitrage. Zarifian et al. (2026) apply random matrix theory to crisis detection.

3 Data and Empirical Design

3.1 Data

I use hourly perpetual futures data for 18 contracts on Binance: BTCUSDT, ETHUSDT, SOLUSDT, BNBUSDT, XRPUSDT, ADAUSDT, DOGEUSDT, LINKUSDT, DOTUSDT, LTCUSDT, AVAXUSDT, UNIUSDT, ATOMUSDT, FILUSDT, NEARUSDT, APTUSDT, ARBUSDT, OPUSDT. Fifteen contracts cover January 2021 to December 2024; three (APTUSDT, ARBUSDT, OPUSDT) cover their full available history. Funding rates are aligned to the hourly grid using the most recent eight-hour observation, never a future one.

3.2 Event definition

An event is an hour in which the standardized funding rate falls below $z < -1.5$ and the previous hour's z was above -0.5 , applied to the 720-hour rolling z -score. The second condition identifies the first hour of each episode. This yields 1,696 non-overlapping events.

3.3 Correction delay

Correction delay is the number of hours from the event until the close price returns to within a fixed percentage of its pre-event 24-hour mean. The baseline threshold is $\pm 1\%$, and the censoring window is 168 hours. Correction is bidirectional.

3.4 Survival analysis and inference

Kaplan-Meier estimates, log-rank tests, and Cox models with vol24 as the covariate. Standard errors are clustered at the year-month level. Cox estimates use Efron's method for ties.

3.5 A coordination-failure framework

Consider N arbitrageurs who each observe a noisy signal $s_i = \theta + \varepsilon_i$ of the true correction demand θ , with $\varepsilon_i \sim N(0, \sigma^2)$. Arbitrageur i acts if $s_i > k$. Correction occurs if at least m of N act. Higher σ reduces the probability of coordination, delaying correction. The model yields three testable predictions: (1) correction delay increases with volatility; (2) the effect concentrates in the early hours; (3) the effect is stronger when volatility is rising.

4 Results

4.1 Descriptive statistics

Table 1: **Table 1. Descriptive statistics.**

Variable	Mean	SD	Median	Min	Max
vol24	0.0103	0.0083	0.0081	0.0009	0.0946
vol168	0.0093	0.0048	0.0083	0.0017	0.0480
vol_ratio	1.0776	0.4280	1.0015	0.1924	2.5560
time_hr	49.72	66.01	11.0	1.0	168.0
event (rate)	0.8013	—	—	0	1

Notes: $n = 1,696$ events across 18 contracts. Overall censoring rate 19.9%.

4.2 Realized volatility and correction speed

High-volatility events take a median of 27.0 hours to correct; low-volatility events take 3.0 hours. The gap is significant (log-rank $\chi^2 = 147.25$, $p < 0.0001$).




Figure 1: **Figure 1. Kaplan-Meier survival curves.**

Table 2: **Table 2. Correction delay by realized volatility.**

Group	n	Median (h)	Censor rate
High volatility	848	27.0	28.3%
Low volatility	848	3.0	11.4%
Gap	—	24.0	16.9 pp

Table 3: **Table 3. Cox model: realized volatility and correction hazard.**

Clustering	HR	SE	95% CI	p
Naive	0.5935	0.0549	[0.5329, 0.6609]	< 0.0001
By contract	0.5935	0.0675	[0.5200, 0.6774]	< 0.0001
By year-month	0.5935	0.0784	[0.5090, 0.6920]	< 0.0001

4.3 Horse race: vol24 versus funding-activity controls

Table 4: **Table 4. Horse race.**

Variable	HR	95% CI (cluster)	Corr. with vol24
vol24 (baseline)	0.589	[0.474, 0.732]	—
vol24 (with controls)	0.584	[0.467, 0.730]	—
abs_mean_24h	1.029	[0.846, 1.253]	0.42
abs_max_24h	0.902	[0.757, 1.073]	0.30
freq_change_24h	1.130	[1.026, 1.244]	−0.02
abs_mean_7d	1.085	[0.853, 1.380]	0.35

4.4 Quintile analysis

Table 5: **Table 5. Correction hazard by volatility quintile.**

Quintile	HR	95% CI (cluster)
Q1 (lowest)	1.000	reference
Q2	0.516	[0.448, 0.594]
Q3	0.790	[0.662, 0.942]
Q4	0.899	[0.757, 1.068]
Q5 (highest)	0.662	[0.554, 0.790]

4.5 Cumulative correction probability

Table 6: **Table 6. Cumulative correction probability.**

Threshold	High P(corr)	Low P(corr)	High P(not)	Low P(not)
$\leq 1\text{h}$	13.65%	43.82%	86.35%	56.18%
$\leq 6\text{h}$	29.65%	57.71%	70.35%	42.29%
$\leq 24\text{h}$	48.94%	73.73%	51.06%	26.27%
$\leq 72\text{h}$	62.24%	83.86%	37.76%	16.14%
$\leq 120\text{h}$	68.47%	86.81%	31.53%	13.19%
$\leq 168\text{h}$	71.80%	88.60%	28.20%	11.40%

4.6 Time-varying effect

Schoenfeld residuals reject the proportional-hazards assumption for vol24 ($\rho = 0.10$, $p = 0.0002$).

I estimate:

$$\lambda(t | x) = \lambda_0(t) \exp(\beta_1 \text{vol24} + \beta_2 \text{vol24} \log(t + 1)),$$

with $\beta_1 = -2.201$ (SE 0.089) and $\beta_2 = 0.3905$ (SE 0.041). The HR is 0.145 at $t = 1\text{h}$, 0.389 at $t = 24\text{h}$, and 0.820 at $t = 168\text{h}$.




Figure 2: **Figure 2. Hazard ratio over time.**

4.7 Volatility gap by calendar year

Table 7: **Table 7. Volatility gap by calendar year.**

Year	n	High median	Low median	χ^2	p
2021	152	70.0	13.0	8.34	0.004
2022	615	24.0	7.0	44.80	< 0.0001
2023	525	41.0	1.0	94.99	< 0.0001
2024	404	20.0	2.0	22.11	< 0.0001

4.8 Volatility-direction sub-samples

Table 8: **Table 8. Volatility-direction sub-samples.**

Sub-sample	n	High median	Low median
Rising volatility	848	55.0	14.0
Falling volatility	848	13.0	1.0

4.9 Correction direction sub-samples

Table 9: **Table 9. Correction direction sub-samples.**

Direction	n	High median	Low median	χ^2	p
Upward (from below)	572	9.0	1.0	38.23	< 0.0001
Downward (from above)	790	12.0	2.0	48.80	< 0.0001
Censored	334	—	—	—	—

4.10 Single-parameter sensitivity

Table 10: **Table 10. Single-parameter sensitivity.**

Scenario	Parameter	High median	Low median	χ^2
Baseline	1.0%, 168h	27.0	3	147.2
Threshold 0.5%	threshold	51	13	99.4
Threshold 2.0%	threshold	11	1	231.9
Window 72h	window	27.0	3	154.9
Window 336h	window	27.0	3	142.2
Split p25	split	19	1	168.9
Split p75	split	34	7	89.9

4.11 Volatility-normalized threshold

Under a volatility-normalized threshold ($\pm \text{vol24} \times \text{pre_mean}$), the high-volatility median is 24.0 hours and the low-volatility median is 12.0 hours ($\log\text{-rank } \chi^2 = 26.41, p < 0.0001$).

4.12 Heterogeneity via recursive partitioning

Table 11: **Table 11. Recursive-partitioning tree leaves.**

Leaf	Rule	n	Mean correction (h)
1	$\text{ratio} \leq 0.678953, \text{vol168} \leq 0.010694$	206	11.53
2	$\text{ratio} \leq 0.678953, \text{vol168} > 0.010694$	80	34.71
3	$0.678953 < \text{ratio} \leq 1.261031, \text{vol168} \leq 0.006376$	261	28.35
4	$0.678953 < \text{ratio} \leq 1.261031, \text{vol168} > 0.006376$	678	49.86
5	$1.261031 < \text{ratio} \leq 2.000576, \text{vol168} \leq 0.010792$	291	65.46
6	$1.261031 < \text{ratio} \leq 2.000576, \text{vol168} > 0.010792$	112	94.45
7	$\text{ratio} > 2.000576$	68	122.62




Figure 3: **Figure 3. Recursive-partitioning tree.**

4.13 Alternative explanations

Contract-specific factors. Within-contract demeaning of vol24 leaves the Cox HR essentially unchanged (0.634 vs baseline 0.594).

Directional factors. Rising: 55 vs 14h ($\chi^2 = 43.61$). Falling: 13 vs 1h ($\chi^2 = 67.60$).

Settlement-cycle timing. All events fall within the same post-settlement window.

Volatility itself. The slow correction is due to the funding-rate extreme itself, rather than to the volatility regime. A placebo test yields 4,186 matched hours. Non-events correct in a median of **10.0 hours**, against **27.0 hours** for events ($\chi^2 = 88.04, p = 6.4 \times 10^{-21}$).




Figure 4: **Figure 4. Placebo test.**

4.14 Censoring robustness

Table 12: **Table 12. Accelerated-failure-time specifications.**

Model	Time ratio	σ	Log-likelihood
Weibull, vol24 winsorized	3.72	2.10	-5928.43
Weibull, log(vol24)	3.03	2.05	-5891.87
Exponential ($\sigma = 1$)	2.32	1.00	-6758.78
Log-normal, vol24 winsorized	3.17	2.24	-5751.86




Figure 5: **Figure 5. Cox-AFT consistency.**

4.15 Natural experiment: the FTX collapse

Table 13: **Table 13. Pre- vs post-FTX.**

Era	n	High median	Low median	Cox HR [95% CI]
Pre-FTX	659	20.0	12.0	0.702 [0.585, 0.841]
Post-FTX	1,037	30.0	1.0	0.570 [0.513, 0.632]

4.16 Policy thresholds

Table 14: **Table 14. Policy thresholds.**

Percentile	vol24 cutoff	Fraction of events	Median correction (h)
50%	0.008141	50.0%	27.0
75%	0.012359	25.0%	34.0
90%	0.018401	10.0%	82.0
95%	0.024073	5.0%	168.0

5 Discussion

5.1 Limits to coordination

There are three orders of coordination failure in funding-rate extremes. First, arbitrageurs must observe the same extreme funding rate. Second, they must agree on the direction of the correction. Third, they must act simultaneously enough to move the price. The evidence in this paper shows that the first two orders are satisfied, but the third fails in high-volatility regimes.

The limits-to-arbitrage literature shows that arbitrageurs fail to correct mispricings when they lack capital (Shleifer and Vishny, 1997) or funding liquidity (Brunnermeier and Pedersen, 2009). The evidence here suggests that arbitrageurs can also fail when they cannot coordinate on a common direction. Three features support this: the effect is present in every contract; it is 4.2 times stronger when volatility is rising; and the placebo test shows that volatility-matched non-event hours correct in a median of 10.0 hours, against 27.0 hours for the events.

5.2 Policy implications

Sometimes a correction is not about solvency. It is about coordination. A venue that wishes to accelerate parity restoration in the tail could apply a multiplier of 2.0 to the funding rate when vol24 exceeds its 90th percentile (0.018401). This targets 10% of events, whose median correction time is 82 hours. The analogy is to the gates-and-fees design used for money market funds. Just as MMFs suspend redemptions during stress to prevent fire sales, a venue could apply a volatility-contingent funding multiplier to slow parity restoration during high-volatility regimes.

5.3 Relation to the literature

The result is consistent with Shleifer and Vishny (1997), Brunnermeier and Pedersen (2009), and Makarov and Schoar (2020). The coordination-failure interpretation connects to Diamond and Dybvig (1983) and Morris and Shin (1998). It is consistent with Pindza (2026), complementary to Willeboordse (2026), and consistent with Zarifian et al. (2026).

5.4 Limitations

L.1 The correction threshold is a reduced-form trigger. **L.2** Events are clustered by contract and time. **L.3** The proportional-hazards assumption is violated for vol24. **L.4** The mechanism is not identified. **L.5** Stress windows cannot be studied at the event level.

6 Conclusion

The evidence introduces a new dimension to the limits-to-arbitrage literature. Capital constraints and funding liquidity have long been recognized as limits to arbitrage. Coordination, this paper shows, is a third. The limits-to-coordination framework introduced in this paper can be applied to other settings where multiple arbitrageurs must coordinate to correct a mispricing, including triangular arbitrage in FX markets, ETF creation-redemption during stress, and Treasury cash-futures basis trades.

A Supplementary Information

Table 15: **Table S1. Intraday distribution.**

UTC hour	n	High median	Low median
00	409	43.0	3.0
01	323	29.0	5.0
08	219	21.0	1.0
09	227	15.0	4.5
16	289	48.0	4.0
17	231	20.5	4.0
Other 18 hours	1	—	—
Total	1,699	—	—

Table 16: **Table S2. Per-contract estimates.**

Contract	n	HR (vol24)	High/Low median
ADAUSDT	90	0.798	22 / 3
APTUSDT	42	0.873	144 / 25
ARBUSDT	38	0.721	22 / 1
ATOMUSDT	106	0.766	23 / 4
AVAXUSDT	99	0.777	25 / 14
BNBUSDT	103	0.301	33 / 2
BTCUSDT	107	0.632	18.5 / 1
DOGEUSDT	109	0.799	29 / 1.5
DOTUSDT	78	0.455	29 / 5
ETHUSDT	99	0.502	15 / 1
FILUSDT	89	0.462	97 / 10.5
LINKUSDT	102	0.742	20 / 6
LTCUSDT	98	0.389	29 / 2
NEARUSDT	113	0.637	21 / 6.5
OPUSDT	66	0.667	24 / 6
SOLUSDT	89	0.490	14 / 5.5
UNIUSDT	142	0.679	19 / 2
XRPUSDT	129	0.672	20 / 1.5

Table 17: **Table S3. Asset-class decomposition.**

Asset class	n	High median	Low median
Major (BTC, ETH)	206	18.0	1.0
Layer 1	617	26.0	5.5
Layer 2	104	21.0	3.0
DeFi	244	20.5	4.0
Payment	227	19.0	2.0
Meme	109	29.0	1.5
Exchange	103	33.0	2.0
Storage	89	97.0	10.5

Table 18: **Table S4. Asset class by year (high/low median, h).**

Class	2021	2022	2023	2024
Major	23/17.5	14/3	15/1	6/1
L1	114/5	14/12	45/3	21/5
L2	NA	18/5.5	25.5/1	24/7
DeFi	13/80	27/7	27/1	25/4
Payment	168/12	18/2.5	22/1	16.5/1
Meme	NA	124.5/37	13.5/1	19/1
Exchange	40.5/35	33/1	19.5/1.5	2/7
Storage	NA	107/9	35/2	142.5/14

Table 19: **Table S5. Out-of-sample and MAD.**

Sample	n	High median	Low median	MAD high/low
In-sample (2021–2023)	1,293	29.0	3.0	28 / 2.0
Out-of-sample (2024)	406	20.0	2.0	19 / 1.0
BTC	107	18.5	1.0	17.5 / 0
ETH	99	15.0	1.0	14.0 / 0
SOL	89	14.0	5.5	13 / 4.5

Table 20: **Table S6. Model evaluation.**

Metric	Value	Benchmark
C-index	0.623	—
Brier (24h)	0.222	0.237

Table 21: **Table S7. Summary of robustness checks.**

Check	Specification	Result
R.1 Baseline	1%, 168h	26.5 / 3, $\chi^2 = 145.9$
R.2 Threshold 0.5%	0.5%	51 / 13
R.3 Threshold 2.0%	2.0%	11 / 1, $\chi^2 = 230.8$
R.4 Window 72h	72h	$\chi^2 = 153.6$
R.5 Window 336h	336h	$\chi^2 = 140.2$
R.6 Grid 4h	4h	28 / 4
R.7 Grid 8h	8h	32 / 8
R.8 Horse race	4 controls	HR=0.584
R.9–R.12 Years	yearly	all significant
R.13 Per-contract Cox	18 contracts	18/18 negative
R.14 Per-contract median	18 contracts	18/18 high>low
R.15 Alternative correction	funding ≥ 0	9 / 8
R.16 Single-parameter	7 scenarios	all significant
R.17–R.18 z -window	360h/1440h	23/3, 27/3
R.19–R.20 z -threshold	−1.0/−2.0	19/2, 38/4
R.21 Large vs small	BTC/ETH/SOL	both directions
R.22 Cumulative	by threshold	gap 22–30 pp
R.23 Intraday	6 settlement hours	6/6 high>low
R.24 Asset class	8 categories	8/8 high>low
R.25 Time halves	2021–22 vs 2023–24	both directions
R.26 Funding magnitude	high vs low	both directions
R.27 Volatility direction	rising vs falling	4.6×
R.28 BNBUSDT removal	drop largest	$\chi^2 = 118.8$
R.29 Quarterly	15 quarters	15/15
R.30 Asset class $\times$ year	29 cells	27/29
R.31 Out-of-sample	2021–23 vs 2024	both directions
R.32 Tail probability	P(>120h)	ratio 2.39
R.33 Monthly consistency	35 months	25/35
R.34 Quarterly consistency	15 quarters	15/15
R.35 Max vol24	SOL 2021-05-19	0.0946
R.36 Economic significance	ratio 8.83×	tail 2.39×
R.37 Correction direction	upward/downward	both directions
R.38 Recursive partitioning	depth 3, min leaf 50	ATE=34.92h
R.39 Vol-normalized threshold	$\pm$ vol24	24 / 12, $\chi^2 = 26.23$

Data availability

The data are openly available in the Binance public data archive at <https://data.binance.vision>. Replication code is available from the author on request.

Author contributions

H.S. conceived and designed the study, acquired and processed the data, implemented the survival analyses, conducted all robustness checks, prepared all figures and tables, and wrote the manuscript.

Declarations

Conflict of interest: The author declares no competing interests.

Use of AI: The author used a large language model for editorial assistance and code generation. All analytical choices, data processing, and conclusions are the author’s own.

References

- [1] Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6), 2201–2238. <https://doi.org/10.1093/rfs/hhn098>
- [2] Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3), 401–419. <https://doi.org/10.1086/261155>
- [3] Epstein, L. G., & Schneider, M. (2003). Recursive multiple-priors. *Journal of Economic Theory*, 113(1), 1–31. [https://doi.org/10.1016/S0022-0531\(03\)00097-8](https://doi.org/10.1016/S0022-0531(03)00097-8)
- [4] Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3), 393–408. <https://www.jstor.org/stable/1805228>
- [5] Hansen, L. P., & Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2), 60–66. <https://doi.org/10.1257/aer.91.2.60>
- [6] Makarov, I., & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2), 293–319. <https://doi.org/10.1016/j.jfineco.2019.07.001>
- [7] Morris, S., & Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3), 587–597. <https://www.jstor.org/stable/116904>
- [8] Pindza, E. (2026). Centralized–decentralized exchange funding rate arbitrage as a basis trade. *Digital Finance*, 8(3). <https://doi.org/10.1007/s42521-026-00213-3>
- [9] Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. *The Journal of Finance*, 52(1), 35–55. <https://doi.org/10.1111/j.1540-6261.1997.tb03807.x>
- [10] Wager, S., & Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523), 1228–1242. <https://doi.org/10.1080/01621459.2017.1319839>
- [11] Willeboordse, F. H. (2026). Fiat versus stablecoin price integration on Kraken. *Digital Finance*, 8(3). <https://doi.org/10.1007/s42521-026-00206-2>
- [12] Zarifian, A., Overbeck, L., Gallus, C., & Blasiak, P. (2026). Crisis detection based on the Marčenko-Pastur law. *Digital Finance*, 8(3). <https://doi.org/10.1007/s42521-026-00214-2>