

Kimchi Premium and Triangular Arbitrage Dislocation Dynamics

Draft v1.3

Songgun Lee
Department of Physics, UC Santa Barbara

As of April 2026

Abstract

This study evaluates minute-level triangular arbitrage dislocations in KRW-linked crypto routing using a 141,200-row export across ADA, DOGE, ETH, SOL, and XRP. Forward and reverse directions are analyzed separately, then aggregated into both a best-per-minute-per-asset series and a global best-per-minute series. Under the Upbit fee schedule, threshold-clearing opportunities are found to be extremely rare and absent under stricter reserved-order and USDT/BTC-equivalent 3-leg friction assumptions. The evidence supports a high short-horizon market-efficiency interpretation for this execution setting.

Contents

1	Research Question	3
2	Executive Summary	3
3	Contributions	3
4	Instrument and Data	3
5	Definitions	4
6	Execution Friction Assumption	4
7	Methodology	5
8	Hypotheses and Test Outcomes	5
9	Summary Statistics	6
9.1	By Asset and Direction	6
9.2	Best-of-Both Aggregates	6
10	Distributional Properties of Global Best-Minute Profit	6
11	Persistence and Short-Lag Dependence	6
12	Direction and Asset-Level Regime Breakdown	7
12.1	Direction-Level	7
12.2	Asset-Level	7

13 Kimchi Premium Bin Sensitivity	7
14 Visualization	8
15 Empirical Finding: High Market Efficiency	8
16 Interpretation	9
17 Fee Scenario Sensitivity	9
18 Limitations	9
19 Next Steps	9
A Reproducibility Appendix	10
A.1 File Lineage	10
A.2 Suggested Regeneration Workflow	10
A.3 Interpretation Guardrails	10

1 Research Question

The central research question is: *Does external macroeconomic dislocation (kimchi-premium volatility) induce exploitable internal microstructural inefficiency within a fiat-isolated exchange order book?*

The empirical test is built around fee-aware tail-event frequency above execution barrier V_0 , with direction-separated and best-direction aggregation layers designed to avoid overstating exploitable edge.

2 Executive Summary

- Dataset scale: 141,200 rows, 5 assets, minute buckets over approximately 10 days.
- Objective: test whether observed dislocations survive realistic 3-leg fee hurdles.
- Main result: profitability tails are thin; positive net-profit minutes are rare even after best-direction selection.
- Fee-consistent result: events above the baseline execution barrier $V_0 = 0.35\%$ are not observed in this sample, and therefore also not observed for stricter barriers.

3 Contributions

Relative to a typical dashboard-style arbitrage note, this draft contributes:

1. A fee-aware execution framing tied directly to the Upbit schedule (KRW general, KRW reserved, USDT/BTC-equivalent 3-leg thresholds).
2. A hierarchical analysis protocol: direction-separated to best-per-asset-per-minute, then global best-per-minute.
3. Explicit negative-result reporting: threshold-clearing rarity is quantified, not inferred.
4. A reproducible paper pipeline (CSV \rightarrow scripts \rightarrow figures/tables \rightarrow LaTeX PDF).

4 Instrument and Data

- Source file: `arb_data_export.csv`
- Time range (UTC): 2026-04-08 10:46:00 to 2026-04-18 07:40:00
- Rows: $N = 141,200$
- Assets: ADA, DOGE, ETH, SOL, XRP
- Directions: Forward / Reverse
- Core fields: `avg_kimchi`, `max_profit`, `min_profit`, `tick_count`

Data Engineering Notes

- `avg_kimchi`: 1-minute arithmetic mean of percentage divergence between Upbit KRW and Binance USDT mid-prices, adjusted by live KRW/USD FX.
- Exchange server timestamps are aligned post-hoc to reduce local hardware/network latency contamination.
- XRP row deficit relative to other assets is treated as likely exchange-side websocket rate-limiting / transient packet loss during high-volume intervals.

5 Definitions

For minute bucket t , asset a , and direction $d \in \{F, R\}$:

$$K_{t,a,d} = \text{avg kimchi premium (\%)} \quad (1)$$

$$\Pi_{t,a,d} = \text{max triangular net profit (\%)} \quad (2)$$

Best-of-both per-minute per-asset:

$$\Pi_{t,a}^* = \max(\Pi_{t,a,F}, \Pi_{t,a,R}) \quad (3)$$

Global best-of-minute:

$$\Pi_t^* = \max_{a,d} \Pi_{t,a,d} \quad (4)$$

Forward route:

$$\text{KRW} \rightarrow \text{BTC} \rightarrow \text{ALT} \rightarrow \text{KRW} \quad (5)$$

Reverse route:

$$\text{KRW} \rightarrow \text{ALT} \rightarrow \text{BTC} \rightarrow \text{KRW} \quad (6)$$

Forward profit mapping used in the edge-node implementation:

$$M = \left(\frac{1}{P_{BTC/KRW}^{ask}} \right) \left(\frac{1}{P_{ALT/BTC}^{ask}} \right) P_{ALT/KRW}^{bid} - 1 \quad (7)$$

Execution mode is pure taker (immediate spread crossing); queue and slippage are excluded from baseline and treated as additional real-world friction.

6 Execution Friction Assumption

Using the Upbit fee guide provided by the author:

- KRW market, general order: Maker/Taker = 0.05%
- KRW market, reserved maker: 0.139%
- BTC/USDT market, general order: Maker/Taker = 0.25%

For a 3-leg triangular execution, the route-level fee barriers are:

$$\tau_{\text{KRW-KRW-BTC}} = 0.05\% + 0.25\% + 0.05\% = 0.35\% \quad (8)$$

$$\tau_{\text{KRW,reserved,all legs}} = 3 \times 0.139\% = 0.417\% \quad (9)$$

$$\tau_{\text{USDT/BTC}} = 3 \times 0.25\% = 0.75\% \quad (10)$$

Primary tail metric in this draft is event frequency above $V_0 = 0.35\%$.

7 Methodology

The analysis is performed in three layers:

1. **Direction-separated layer:** evaluate $\Pi_{t,a,F}$ and $\Pi_{t,a,R}$ independently to capture directional asymmetry.
2. **Asset-level best-direction layer:** compute $\Pi_{t,a}^*$ to represent best feasible routing direction per minute for each asset.
3. **Global best-minute layer:** compute Π_t^* as an optimistic upper envelope across all assets and directions.

This structure intentionally tests both realistic directional execution and an optimistic upper-bound case, reducing the chance of overstating tradable edge.

8 Hypotheses and Test Outcomes

We define three empirical hypotheses for this sample window:

- **H1 (Direction asymmetry):** Forward and reverse routes differ materially in realized edge distribution.
- **H2 (Executable-edge rarity):** Even optimistic best-minute aggregation rarely clears realistic fee thresholds.
- **H3 (Kimchi monotonicity):** Higher kimchi-premium regimes do not guarantee higher fee-cleared opportunity rates.

Observed outcomes in this sample:

- H1 supported: forward mean exceeds reverse mean (less negative) and has higher positive-event frequency.
- H2 strongly supported: zero exceedance at $V_0 = 0.35\%$, with zero exceedance also observed at 0.417% and 0.75% for global best-minute.
- H3 supported: bin-wise results show no robust monotone increase in fee-cleared rates as kimchi premium rises.

9 Summary Statistics

9.1 By Asset and Direction

Table 1: Direction-separated profitability summary

Asset	Dir	N	Mean	Std	Min	Max	$P(\Pi > 0)$	$P(\Pi > V_0)$
ADA	F	14210	-1.3877	0.6441	-4.0253	0.1522	0.000633	0.000000
ADA	R	14210	-1.5917	0.7930	-4.5642	-0.1922	0.000000	0.000000
DOGE	F	14210	-1.4591	0.3912	-2.7039	-0.1189	0.000000	0.000000
DOGE	R	14210	-1.6219	0.3200	-2.3922	-0.3843	0.000000	0.000000
ETH	F	14210	-0.6900	0.1377	-1.1212	0.1256	0.000422	0.000000
ETH	R	14210	-0.7040	0.1125	-0.9426	0.1380	0.000141	0.000000
SOL	F	14210	-0.8059	0.2668	-2.0355	0.0287	0.000211	0.000000
SOL	R	14210	-0.7247	0.1567	-1.1677	0.0207	0.000070	0.000000
XRP	F	14210	-0.4987	0.2191	-1.2628	0.0399	0.001267	0.000000
XRP	R	13310	-0.6181	0.1873	-0.9696	0.2545	0.000075	0.000000

9.2 Best-of-Both Aggregates

Table 2: Best-of-both and global-best minute aggregates

Series	N	Mean	Std	Min	Max	$P(\Pi > 0)$	$P(\Pi > V_0)$
$\Pi_{t,a}^*$ (best per-minute per-asset)	71050	-0.8122	0.4171	-2.4768	0.2545	0.000563	0.000000
Π_t^* (global best per-minute)	14210	-0.3654	0.1561	-0.7986	0.2545	0.002815	0.000000

10 Distributional Properties of Global Best-Minute Profit

Let Π_t^* denote the global best-per-minute series. Empirical moments for this sample:

- Mean: -0.3654%
- Standard deviation: 0.1561%
- Skewness: -0.2420
- Quantiles: $q_{0.05} = -0.6521\%$, $q_{0.50} = -0.3531\%$, $q_{0.95} = -0.1212\%$

All central quantiles remain negative, reinforcing the high-efficiency interpretation under this execution model.

11 Persistence and Short-Lag Dependence

The lag-1 autocorrelation of Π_t^* is:

$$\text{ACF}(1) \approx 0.9332 \quad (11)$$

This indicates strong minute-to-minute persistence in the dislocation state. Given high-frequency overlap and market microstructure continuity, this should be interpreted as state persistence rather than independent repeated arbitrage opportunity.

12 Direction and Asset-Level Regime Breakdown

12.1 Direction-Level

Table 3: Forward vs reverse direction statistics

Direction	N	Mean max_profit (%)	$P(\Pi > 0)$
Forward	71050	-0.9683	0.000507
Reverse	70150	-1.0576	0.000057

12.2 Asset-Level

Table 4: Asset-level profitability summary

Asset	N	Mean max_profit (%)	$P(\Pi > 0)$
ADA	28420	-1.4897	0.000317
DOGE	28420	-1.5405	0.000000
ETH	28420	-0.6970	0.000281
SOL	28420	-0.7653	0.000141
XRP	27520	-0.5564	0.000690

13 Kimchi Premium Bin Sensitivity

Using global best-minute observations, kimchi-premium bins are defined and profitability frequencies are computed.

Table 5: Global best-minute by kimchi-premium regime

Kimchi bin (%)	N	Mean (%)	$P(\Pi > 0)$	$P(\Pi > 0.35\%)$
< 0	2965	-0.3873	0.007083	0.000000
$[0, 0.25)$	3087	-0.3666	0.001620	0.000000
$[0.25, 0.5)$	4016	-0.3692	0.001743	0.000000
$[0.5, 0.75)$	3147	-0.3497	0.002224	0.000000
≥ 0.75	995	-0.3305	0.000000	0.000000

No monotone “higher-kimchi, higher-executable-edge” relation appears in this short sample after fee-aware filtering.

14 Visualization

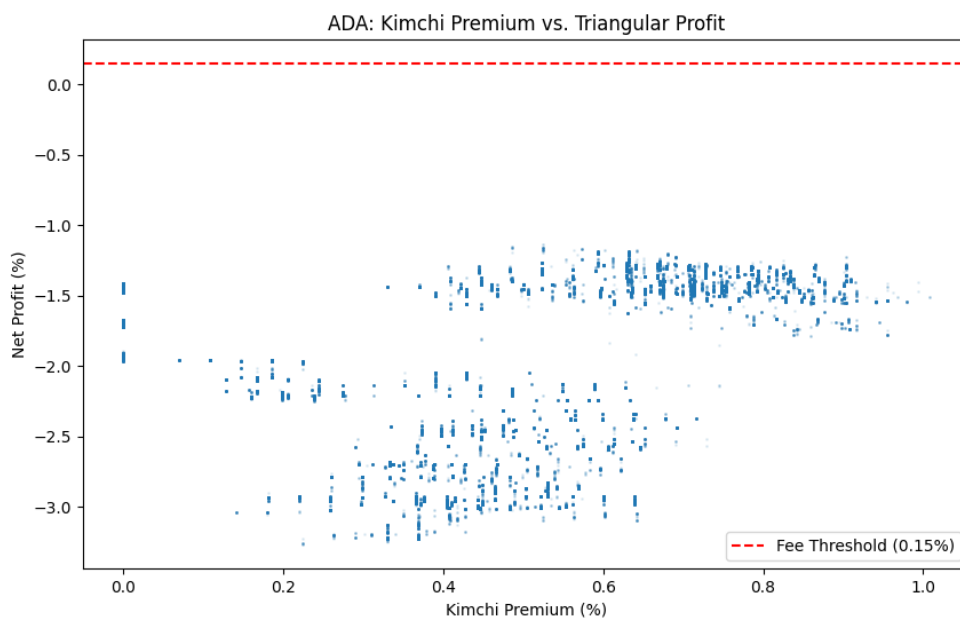


Figure 1: ADA compressed-history view: Kimchi premium vs. triangular net profit.

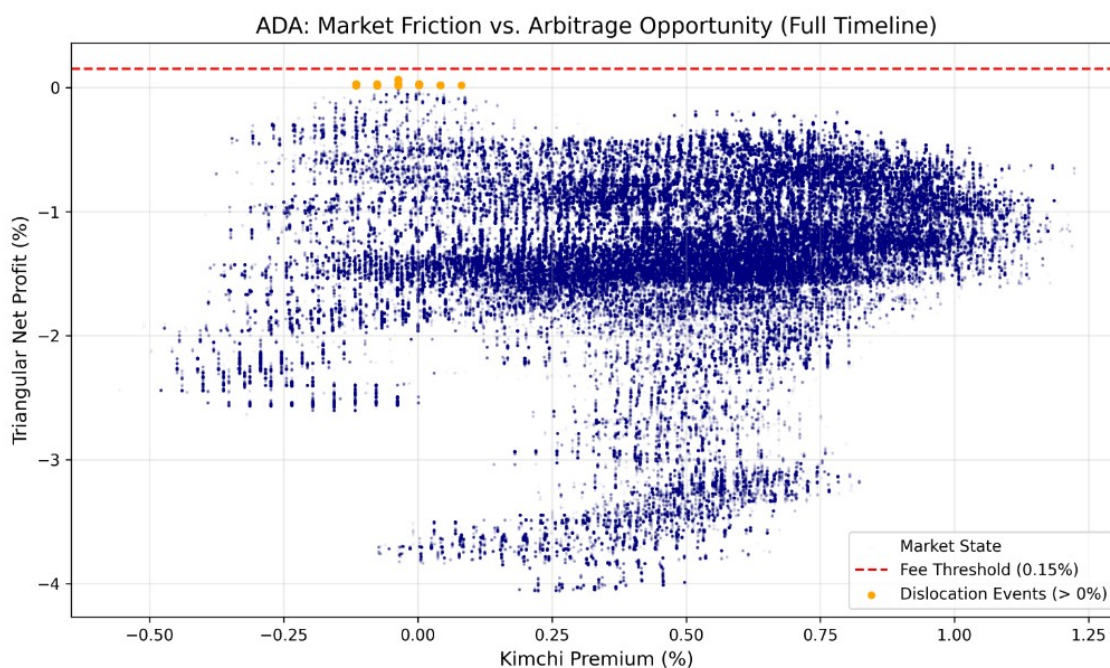


Figure 2: ADA full timeline view with fee threshold and dislocation events.

15 Empirical Finding: High Market Efficiency

The dominant empirical result is strong execution inefficiency for arbitrage capture (i.e., market *efficiency* against exploitable edge):

- Nearly all observations remain below zero net profit after modeled friction.
- Even under best-of-both direction selection, positive-minute probability remains very low.
- Threshold-clearing events above the baseline barrier $V_0 = 0.35\%$ are absent in the observed window.

For ADA best-of-both minutes, the Kimchi-premium-to-profit linear correlation is:

$$\text{Corr}(K, \Pi^*) \approx 0.1735 \quad (12)$$

which indicates only weak monotone linkage in this sample.

16 Interpretation

Within this observed period, cross-venue dislocations are mostly absorbed before overcoming execution friction. In practical terms, the market behaves as highly efficient at the minute-level for this strategy specification.

17 Fee Scenario Sensitivity

For global best-per-minute profitability Π_t^* :

- $P(\Pi_t^* > 0.15\%) = 0.000141$
- $P(\Pi_t^* > 0.35\%) = 0$
- $P(\Pi_t^* > 0.417\%) = 0$
- $P(\Pi_t^* > 0.75\%) = 0$

In this sample, opportunities that clear the baseline mixed-leg barrier (0.35%), the KRW reserved-order barrier (0.417%), or the USDT/BTC-equivalent barrier (0.75%) are not observed.

18 Limitations

- Sample spans roughly 10 days, so regime-level conclusions are preliminary.
- Baseline excludes slippage and queue dynamics; practical execution may require $M > 0.40\%$ to survive additional frictions.
- Direction-level asymmetries may be influenced by routing constraints not explicitly modeled in this draft.

19 Next Steps

1. Extend to multi-month / multi-regime dataset.
2. Add exchange-specific fee + slippage decomposition per route (leg-by-leg).
3. Build conditional models $P(\Pi > \tau \mid K, \text{tick_count}, a, d)$.
4. Add event overlays (macro releases, local KRW liquidity shocks).

A Reproducibility Appendix

A.1 File Lineage

- Input data: `arb_data_export.csv`
- Plot scripts: `generate_plot.py`
- Paper generator: `generate_paper.py`
- White paper source: `crypto_kimchi_draft_v1.tex`
- Output figures: `ada_analysis.png`, `ada_full_analysis.png`
- Output PDF: `crypto_kimchi_draft_v1.pdf`

A.2 Suggested Regeneration Workflow

```
# 1) (Optional) regenerate the dynamic paper source from DB export logic
python3 generate_paper.py
```

```
# 2) regenerate plots
python3 generate_plot.py
```

```
# 3) compile white paper draft (run twice for TOC)
pdflatex -interaction=nonstopmode -halt-on-error crypto_kimchi_draft_v1.tex
pdflatex -interaction=nonstopmode -halt-on-error crypto_kimchi_draft_v1.tex
```

A.3 Interpretation Guardrails

- This is a short-window empirical draft, not a full-cycle market invariance claim.
- Threshold exceedance depends on route design, queue priority, and unmodeled slippage.
- Negative findings remain informative: they constrain where exploitable edge is *not* present.