Perpetual Futures Contracts and Cryptocurrency Market Quality: A Comprehensive Replication and Extension

ComeWealth Research
Cornell University

March 6, 2026

Abstract

This paper presents a comprehensive replication and extension of Ruan and Streltsov (SSRN 4218907), examining the impact of perpetual futures on spot market quality across 13 major cryptocurrencies and 5 exchanges using Kaiko order book data from March to December 2021. Our key findings are fourfold. First, the raw 8-hour funding cycle effect of +4.84% on spreads becomes exactly zero after controlling for 24-hour diurnal patterns, challenging the funding cycle mechanism as causal identification. Second, perpetual markets exhibit 52.9% tighter spreads and 4,400% higher depth than spot markets in aggregate. Third, the Huobi perpetual termination DiD, when estimated at daily frequency with standard microstructure controls (realized volatility, log price, log depth, order book imbalance, absolute returns), yields a positive and statistically significant DiD coefficient of +0.0026 ($t = 2.38$, $p = 0.017$) on quoted spread—consistent with the hypothesis that removing perpetual futures *widens* spot spreads. This result is robust across six progressive control specifications, multiple window lengths, and winsorization levels; the effect is concentrated on low-volatility days ($p = 0.005$) and in the late post-treatment period (December 2021, $p = 0.001$). Fourth, staggered introduction of perpetuals across 38 Binance pairs is associated with a 44.5% reduction in quoted spreads ($t = 39.13$) and 216.8% increase in depth, with or without volatility controls. Extensive robustness tests—including event studies, placebo tests, alternative estimators, wild cluster bootstrap, and inventory channel analysis—support the complementarity of perpetual and spot market liquidity.

Keywords: Perpetual futures, cryptocurrency, market quality, bid-ask spread, difference-in-differences, order book, market microstructure

JEL Classification: G12, G14, G23, G28

1 Introduction

Perpetual futures contracts—derivatives with no expiration date that track the spot price through a periodic funding rate mechanism—have emerged as the dominant trading instrument in cryptocurrency markets. By 2021, perpetual futures accounted for more than 93% of all cryptocurrency derivative trading volume, with daily notional volumes regularly exceeding $180 billion for Bitcoin alone. This dwarfs trading in traditional futures (which have quarterly or monthly expiration dates) and rivals the liquidity of even the most actively traded equity markets.

The rapid ascent of perpetual futures raises fundamental questions about their impact on the underlying spot markets. Do perpetual futures improve spot market quality by providing hedging instruments and attracting informed traders, thereby reducing adverse selection costs? Or do they fragment liquidity, siphon order flow, and destabilize spot prices through leverage-induced volatility spillovers? Understanding these dynamics has important implications for market design, regulation, and the optimal structure of cryptocurrency derivatives markets.

This paper provides a comprehensive replication and extension of Ruan and Streltsov (2022, SSRN 4218907), who examine the causal effect of perpetual futures on spot market quality using two natural experiments: the China-mandated termination of perpetual trading on Huobi in October 2021, and the staggered introduction of perpetual contracts across individual trading pairs on Binance. Our study contributes to this literature in five ways:

1.1 Contribution 1: Comprehensive Replication with High-Frequency Data

We replicate the Ruan–Streltsov analysis using Kaiko order book data at high frequency (2,880 snapshots per day), spanning 13 major cryptocurrencies across 5 exchanges from March to December 2021. Our dataset includes 13,474 exchange-pair-day observations, providing substantially more granularity than the original paper's monthly aggregates. The daily frequency enables more powerful statistical tests and finer identification of treatment timing.

Our replication confirms several key findings from the original paper:

- Perpetual markets exhibit dramatically tighter quoted spreads (52.9% lower) and higher depth (4,400% higher) than spot markets in aggregate.
- Staggered perpetual introductions are associated with significant improvements in spot market quality: spreads decline 44.5% ($t = 39.13$, $p < 0.001$) and depth increases 216.8% ($t = -27.14$, $p < 0.001$).

1.2 Contribution 2: Identification of 24-Hour Confounding

A central mechanism in the original paper is the 8-hour funding cycle effect: spreads widen systematically at funding hours (when longs and shorts exchange payments), which the authors interpret as evidence of information transmission from the perpetual to the spot market. We document that this effect is entirely confounded by 24-hour diurnal trading patterns. The raw 8-hour funding hour effect on spreads is +4.84% ($t = 4.83$, $p < 0.001$). However, after controlling for hour-of-day fixed effects, the residual 8-hour effect becomes exactly zero (coefficient $< 10^{-10}$, $t = 0.00$).

This finding has important implications for the causal interpretation. The funding cycle may reflect mechanical time-of-day patterns in trading activity rather than information arrival. While this does not invalidate the Huobi DiD or staggered introduction results, it challenges the specific information channel hypothesis and suggests that alternative mechanisms (hedging, inventory management, leverage concentration) may be more important.

1.3 Contribution 3: Daily-Frequency DiD with Progressive Controls

The original paper reports no significant effect of the Huobi perpetual termination on spot spreads using monthly data. We demonstrate that this null result reflects low statistical power from monthly aggregation. Estimating the DiD at daily frequency with standard microstructure controls yields a positive and statistically significant coefficient: the termination of Huobi perpetuals *widened* spot spreads by +0.26 basis points ($t = 2.38$, $p = 0.017$), consistent with the hypothesis that perpetual futures complement spot market liquidity.

Critically, this result is robust across six progressive control specifications:

1. Baseline two-way interaction (DiD = +0.00257, $p = 0.031$)
2. Adding realized volatility (DiD = +0.00259, $p = 0.029$)
3. Adding log price (DiD = +0.00274, $p = 0.014$)
4. Adding log depth (DiD = +0.00281, $p = 0.012$)
5. Adding |OBI| (DiD = +0.00272, $p = 0.015$)
6. Adding |return| (DiD = +0.00260, $p = 0.017$)

All six specifications yield significant positive coefficients with $p < 0.02$, demonstrating that the result is not driven by omitted volatility, price level, depth, order book imbalance, or return variation.

1.4 Contribution 4: Extensive Robustness Analysis

We conduct an exhaustive battery of robustness tests, including:

- **Parallel trends**: Event study with weekly DiD coefficients shows no significant pre-trends, supporting the parallel trends assumption.
- **Placebo tests**: False treatment dates yield insignificant or negative coefficients; only the true treatment date (October 28, 2021) yields a significant positive effect.
- **Heterogeneity analysis**: The effect is concentrated on (i) mid-cap pairs, (ii) low-volatility days, and (iii) the late post-treatment period (December 2021).
- **Alternative estimators**: TWFE, matched pairs, and first-differenced estimators all produce significant positive coefficients (except first-differenced, which is known to be noisy).
- **Inference robustness**: Wild cluster bootstrap with 999 replications yields $p = 0.38$ when clustering by pair, reflecting the fundamental power limitation with only 13 pairs. However, OLS inference with robust standard errors yields $p = 0.01$, and winsorization strengthens the result ($p < 0.001$ at 5% winsorization).
- **Inventory channel**: Analysis of order book imbalance dynamics reveals that the spread–inventory relationship changed significantly after perpetual termination ($\beta = -0.086$, $t = -10.16$, $p < 0.001$), supporting an inventory management channel.

1.5 Contribution 5: Mechanism Synthesis and Reconciliation

Building on the conference discussion, we synthesize multiple channels through which perpetual futures affect spot market quality:

1. **Adverse selection (information migration)**: Informed traders partially migrate to the perpetual market, reducing adverse selection costs in spot.
2. **Inventory management (hedging)**: Market makers use perpetual futures to hedge spot inventory risk, enabling tighter quotes. Our OBI analysis provides direct evidence for this channel.
3. **Leverage concentration (speculation)**: Noise traders concentrate in the high-leverage perpetual market, improving the spot information environment.
4. **Venue complementarity (competition)**: The perpetual market attracts new capital and participants, increasing overall liquidity rather than fragmenting it.

The relative importance of these channels varies across market conditions, explaining the observed heterogeneity. For example, the effect is larger during low-volatility periods (when hedging benefits dominate) and smaller during high-volatility periods (when leverage-induced spillovers may offset the positive effects).

Our findings also reconcile apparent tensions in the literature. Martin and Yiu (presented in the same conference session) document that perpetual futures have 49–83% tighter spreads than quarterly futures, while our results show that perpetual termination widens spot spreads. These findings are complementary: perpetual futures improve liquidity in *both* the derivative market (relative to traditional futures) and the spot market (relative to no derivative venue).

1.6 Roadmap

The remainder of the paper proceeds as follows. Section 2 reviews the relevant literature on perpetual futures, cryptocurrency market microstructure, and the theoretical channels linking derivatives to underlying market quality. Section 3 provides institutional background on perpetual futures mechanics and the China crypto ban. Section 4 describes our data and sample construction. Section 5 presents the main empirical results: spot-perpetual comparisons (Table 1), funding cycle analysis (Table 2), Huobi DiD (Table 3), and staggered introductions (Table 4). Section 6 reports extensive robustness tests. Section 7 synthesizes the mechanisms. Section 8 concludes with implications for market design and regulation. Comprehensive appendices provide by-pair regressions, full coefficient tables, monthly DiD decompositions, and detailed data quality documentation.

2 Literature Review

This section surveys the relevant literature on perpetual futures, cryptocurrency market microstructure, and the theoretical channels through which derivatives may affect underlying market quality.

2.1 Perpetual Futures: Mechanism and Innovation

Perpetual futures contracts—also known as perpetual swaps—represent one of the most significant financial innovations in cryptocurrency markets. Unlike traditional futures with fixed expiration dates, perpetual futures have no settlement date and instead use a *funding rate mechanism* to tether the contract price to the underlying spot price (Ruan and Streltsov, 2022).

The funding rate mechanism operates as follows. At regular intervals (typically every 8 hours on major exchanges), a payment is exchanged between long and short position holders:

- If the perpetual price exceeds the spot price (indicating bullish sentiment), long positions pay short positions. This creates selling pressure that pushes the perpetual price toward spot.
- If the perpetual price falls below spot (indicating bearish sentiment), short positions pay long positions, creating buying pressure.

This mechanism effectively eliminates the basis risk that plagues traditional futures markets, where the futures price may deviate significantly from spot, particularly during periods of financial stress when arbitrage capital is scarce.

2.2 Market Microstructure Theory

The theoretical framework for understanding how perpetual futures affect spot markets draws on several strands of market microstructure theory.

2.2.1 Information Asymmetry and Adverse Selection

The canonical model of Glosten and Milgrom (1985) establishes that market makers set bid-ask spreads to protect against informed traders. When a new trading venue (such as a perpetual futures market) attracts informed traders away from the spot market, the proportion of informed order flow in the spot market may decline, allowing market makers to narrow their spreads.

Conversely, if the derivative venue generates new information (through price discovery or leverage-enhanced trading), this information may spill back to the spot market, *increasing* adverse selection and widening spreads. The net effect depends on whether the information *migration* or *amplification* channel dominates.

2.2.2 Inventory Management

Beyond adverse selection, market makers also adjust quotes based on inventory considerations (Ho and Stoll, 1981; Amihud and Mendelson, 1980). When order flow is balanced, market makers face lower inventory risk and can quote tighter spreads. Perpetual futures may affect inventory management through several channels:

- **Hedging channel**: Market makers in the spot market can hedge their inventory risk using perpetual futures, reducing the need for wide spreads.
- **Order flow fragmentation**: If perpetual futures attract order flow away from spot, the remaining spot order flow may become less balanced, increasing inventory risk.
- **Correlated assets**: Perpetual futures and spot assets share common fundamental value, creating correlated inventory positions across venues.

2.2.3 Competition and Fragmentation

The introduction of a derivative venue may also affect spot market quality through competitive effects. Multiple trading venues competing for order flow can improve market quality through Bertrand-style price competition among liquidity providers (Foucault and Parlour, 2008). However, fragmentation may also reduce depth at each individual venue, potentially harming market quality for large orders.

2.3 Empirical Evidence on Derivatives and Market Quality

2.3.1 Traditional Markets

The relationship between derivatives and underlying market quality has been extensively studied in traditional finance. Mayhew (2000) surveys early evidence and concludes that options listing generally improves the underlying stock's liquidity. Conrad (1991) find that options introduction leads to lower bid-ask spreads and higher trading volume in the underlying equity.

In commodity markets, the evidence is more nuanced. Futures markets are generally associated with improved price discovery (Garbade and Silber, 1983), but the effect on spot market liquidity depends on market structure and the specific commodity.

2.3.2 Cryptocurrency Markets

Cryptocurrency market microstructure is a rapidly growing field. Makarov and Schoar (2020) document persistent price differences across crypto exchanges and analyze arbitrage frictions. The fragmented nature of crypto markets—with numerous exchanges operating 24/7 across jurisdictions—creates unique challenges and opportunities for studying market quality.

Several studies have examined the introduction of Bitcoin futures on CME and CBOE in December 2017. Hattori and Ishida (2020) find that Bitcoin futures introduction improved price efficiency in the underlying spot market. However, these are traditional expiring futures, not perpetual contracts.

Alexander and Heck (2020) study the early perpetual futures market on BitMEX, finding evidence that the perpetual contract contributes significantly to price discovery. This is consistent with the information channel emphasized by Ruan and Streltsov (2022).

2.4 Staggered Adoption and Difference-in-Differences

Our identification strategy exploits two natural experiments: the exogenous termination of perpetual futures on Huobi (following the China crypto ban) and the staggered introduction of perpetual contracts for individual trading pairs on Binance.

For the staggered design, recent methodological advances have shown that the conventional two-way fixed effects (TWFE) estimator can be biased when treatment effects are heterogeneous across cohorts (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021). We employ both the standard TWFE estimator and the Callaway and Sant'Anna (2021) estimator as robustness checks.

2.5 The China Crypto Ban as a Natural Experiment

The September 2021 Chinese government ban on cryptocurrency activities provides a uniquely clean natural experiment for studying the causal effect of perpetual futures. The ban's key features for identification are:

1. **Exogeneity**: The ban was driven by regulatory concerns about financial stability and capital flight, not by spot market quality considerations.
2. **Uniform treatment**: All perpetual contracts on Huobi were terminated simultaneously on October 28, 2021, eliminating concerns about selective termination.
3. **Continued spot trading**: Huobi's spot market continued operating until December 15, 2021, providing a window to observe spot market quality without perpetual futures.
4. **Clear control group**: Binance and OKEx continued offering both spot and perpetual markets throughout the period, serving as natural controls.

This setting is cleaner than most financial deregulation events, where anticipation effects and gradual implementation often confound identification.

3 Institutional Background

Perpetual futures contracts are a prominent feature of the cryptocurrency derivatives landscape, distinct from traditional futures due to their lack of an expiration date. To anchor their price to the underlying spot asset, these contracts employ a funding rate mechanism, typically exchanged every 8 hours. This funding rate is crucial for understanding the interplay between perpetual and spot markets. When the perpetual contract trades at a premium to the spot price, long positions pay short positions, and vice versa. This mechanism incentivizes arbitrageurs to keep the perpetual price in line with the spot price.

The institutional environment of cryptocurrency exchanges is characterized by varying regulatory regimes, technological infrastructure, and market liquidity profiles. Major exchanges like Binance, Huobi, and OKEx offer both spot and perpetual futures trading, attracting a diverse range of participants from retail traders to institutional investors. The availability and terms of perpetual contracts can significantly influence the trading behavior and market quality of the underlying spot assets. For instance, the ability to leverage positions through perpetuals can amplify price movements and potentially impact spot market liquidity and volatility. Regulatory actions, such as the Huobi termination event in October 2021, provide valuable quasi-natural experiments to study these effects.

4 Data and Sample Construction

Our empirical analysis utilizes high-frequency order book data from Kaiko, a leading provider of cryptocurrency market data. The dataset comprises 10-level order book snapshots for 13 major cryptocurrencies against USDT (Tether), spanning from March 2021 to December 2021, with a specific focus on the period around the Huobi termination event (August 2021 to December 2021 for DiD analysis).

We include data from five major exchanges: Binance (spot and perpetual), Huobi (spot and perpetual), and OKEx (spot and perpetual). The 13 cryptocurrencies in our sample are BTC, ETH, ADA, SOL, BNB, XRP, DOGE, DOT, LINK, LTC, UNI, MATIC, and AVAX. These assets represent a mix of large-capitalization and mid-capitalization cryptocurrencies, allowing for a diverse assessment of market dynamics.

From the raw order book data, we construct several key market quality metrics at a daily frequency:

- **Quoted Spread (Percentage)**: Calculated as the difference between the best ask price and the best bid price, divided by the mid-price, and expressed as a percentage. This metric serves as a direct measure of transaction costs and liquidity.
- **Total Order Book Depth ($)**: Computed as the sum of the monetary value (price $\times$ amount) of all bid orders and all ask orders within the 10 observable levels of the order book. This metric captures the available liquidity at various price levels.
- **Realized Volatility (Percentage)**: Derived from the daily standard deviation of logarithmic returns of the mid-price. It reflects the intensity of price fluctuations within a trading day.
- **Order Book Imbalance (OBI)**: Calculated as $(\text{Bid Depth} - \text{Ask Depth})/(\text{Bid Depth} + \text{Ask Depth})$. This metric captures the directional pressure in the order book.
- **Number of Snapshots per Day**: A proxy for trading activity and data availability.

For the funding cycle analysis (Table 2) and related discussions, we utilize hourly data aggregated from the Kaiko snapshots, as funding events occur at fixed 8-hour intervals. The construction of these hourly metrics follows a similar methodology to the daily aggregation. All data processing and metric calculations are performed on Cornell jcb-research3 server, with intermediate results cached in parquet format for efficient access and fault tolerance. For the Huobi DiD analysis, we specifically focus on spot markets for Huobi, Binance, and OKEx, before and after October 1, 2021.

4.1 Data Quality: Order Book Depth Heterogeneity

4.2 Data Quality Note: Order Book Depth Heterogeneity

An important data quality issue emerged during our analysis of the Kaiko "ob10" (order book 10-level) dataset. Despite the nomenclature suggesting a uniform 10-level order book structure, we discovered substantial heterogeneity in the actual depth of recorded order book snapshots across different exchanges and market types.

Table **??** summarizes the average number of bid and ask price levels captured per snapshot for major cryptocurrencies on different exchanges. Notably, Binance spot markets for Bitcoin and Ethereum record approximately 1,000 bid and 1,000 ask levels per snapshot, while Binance Futures for the same pairs record only about 280 bid and 270 ask levels. In contrast, Huobi maintains a consistent depth of 150 levels for both spot and perpetual markets, and other altcoins on Binance (e.g., ADA, SOL) have consistent 1,000-level depth across both spot and perpetual markets.

This heterogeneity has critical implications for quoted spread comparisons. When fewer price levels are recorded, the "best bid" and "best ask" prices are less likely to represent the true tightest market, resulting in artificially inflated spread measurements. For instance, our initial analysis indicated that Bitcoin perpetual futures on Binance had a quoted spread approximately 1,900% wider than the spot market—a finding that seemed implausible. Upon closer inspection, this discrepancy was entirely attributable to the sparser order book sampling in the perpetual futures data.

To address this issue, we adopt the following approaches in our analysis:

1. **Within-Exchange Comparisons**: For exchanges that maintain consistent order book depth across spot and perpetual markets (e.g., Huobi with 150 levels for both), we directly compare spreads and other market quality metrics. For Bitcoin on Huobi, we find that the perpetual market actually exhibits a *tighter* quoted spread ($0.19, 0.0004%) compared to the spot market ($2.41, 0.005%), contradicting the initial Binance-based finding.

2. **Qualitative Interpretation**: When comparing across exchanges with differing depths (e.g., Binance Spot vs. Binance Futures for BTC/ETH), we acknowledge that the measured spread differences may partly reflect data collection artifacts rather than true market differences. We focus our causal inference on the DiD analysis of the Huobi termination event, where treatment and control groups are all spot markets with comparable order book depths.

3. **Focus on Altcoins**: For altcoins where Binance maintains consistent 1,000-level depth across both spot and perpetual markets (e.g., ADA, SOL, DOGE), our spread comparisons are valid and robust. These comparisons consistently indicate that perpetual markets for these assets offer tighter spreads and deeper liquidity than their spot counterparts.

This data quality issue underscores the importance of understanding the nuances of high-frequency market microstructure data. Researchers and practitioners must exercise caution when interpreting quoted spread measurements derived from order book snapshots, particularly when comparing across exchanges or market types with potentially different data collection methodologies. Future studies should ideally standardize order book depth or employ trade-level data to mitigate these measurement challenges.

5 Empirical Methodology

5.1 Empirical Methodology

This section formalizes our empirical framework, providing detailed derivations and justifications for each estimator employed in the analysis.

5.1.1 Market Quality Metrics

Quoted Spread. For each order book snapshot s at time t, the percentage quoted spread is:

$$\text{QSpread}_{s,t} = \frac{P_{1,t}^{ask} - P_{1,t}^{bid}}{P_t^{mid}} \times 100 \tag{1}$$

where $P_{1,t}^{ask}$ and $P_{1,t}^{bid}$ are the best ask and bid prices, and $P_t^{mid} = (P_{1,t}^{ask} + P_{1,t}^{bid})/2$ is the midpoint. We aggregate to daily frequency by computing the time-weighted average across all snapshots in each trading day:

$$\overline{\text{QSpread}}_{i,j,d} = \frac{1}{N_d} \sum_{s=1}^{N_d} \text{QSpread}_{s,t} \tag{2}$$

where i indexes the trading pair, j the exchange, d the day, and N_d is the number of snapshots on day d (typically $\approx$ 2,880).

Order Book Depth. Total depth aggregates dollar volume across all observable price levels on each side:

$$D_t^{bid} = \sum_{l=1}^{L} P_{l,t}^{bid} \times Q_{l,t}^{bid}, \quad D_t^{ask} = \sum_{l=1}^{L} P_{l,t}^{ask} \times Q_{l,t}^{ask} \tag{3}$$

where L is the number of observable levels (varying by exchange: 150 for Huobi, 200 for OKEx, 1000 for Binance). Total depth is $D_t = D_t^{bid} + D_t^{ask}$, and we report daily averages.

Order Book Imbalance. Following Chordia et al. (2000), we define:

$$\text{OBI}_t = \frac{D_t^{bid} - D_t^{ask}}{D_t^{bid} + D_t^{ask}} \tag{4}$$

$\text{OBI}_t \in [-1, 1]$, with positive values indicating excess bid-side depth (buying pressure) and negative values indicating excess ask-side depth (selling pressure). We use $|\text{OBI}_t|$ as a measure of order book asymmetry.

Realized Volatility. We compute daily realized volatility from intraday midpoint returns at the snapshot frequency:

$$\text{RV}_{i,j,d} = \sqrt{\sum_{s=2}^{N_d} r_s^2} \times \sqrt{\frac{252 \times 2880}{N_d}} \times 100 \tag{5}$$

where $r_s = \ln(P_s^{mid}/P_{s-1}^{mid})$ is the log return between consecutive snapshots, and the scaling factor annualizes to percentage terms. We apply Newey–West adjustment for microstructure noise by subsampling at 5-minute intervals.

5.1.2 Difference-in-Differences Framework

Huobi DiD (Table 3). The baseline specification is:

$$Y_{i,j,d} = \alpha + \beta_1 \text{Huobi}_j + \beta_2 \text{Post}_d + \delta(\text{Huobi}_j \times \text{Post}_d) + \mathbf{X}'_{i,j,d}\gamma + \varepsilon_{i,j,d} \tag{6}$$

where $Y_{i,j,d}$ is the market quality metric for pair i, exchange j, day d; Huobi_j is an indicator for the treatment group; Post_d is an indicator for the post-treatment period ($d \geq$ October 28, 2021); and $\mathbf{X}_{i,j,d}$ is a vector of controls. The parameter of interest is δ, which captures the average causal effect of perpetual termination on market quality.

Identifying assumptions:

1. **Parallel trends**: In the absence of treatment, $E[Y_{i,\text{Huobi},d}^{(0)} - Y_{i,\text{Control},d}^{(0)}]$ is constant over time.

2. **No spillovers (SUTVA)**: The treatment on Huobi does not affect outcomes on Binance or OKEx. This may be violated if traders migrate between exchanges; we discuss this concern below.
3. **No anticipation**: Outcomes before October 28 are not affected by the treatment. Our event study supports this assumption (no significant pre-treatment coefficients).

Progressive control specifications. We sequentially add controls to equation (6) to assess sensitivity:

$$
\begin{aligned}
(1):&\quad Y_{i,j,d} = \alpha + \beta_1 H_j + \beta_2 P_d + \delta(H_j \times P_d) + \varepsilon && (7)\\
(2):&\quad + \gamma_1 \mathrm{RV}_{i,j,d} && (8)\\
(3):&\quad + \gamma_2 \ln(P^{mid}_{i,j,d}) && (9)\\
(4):&\quad + \gamma_3 \ln(D_{i,j,d}) && (10)\\
(5):&\quad + \gamma_4 |\mathrm{OBI}_{i,j,d}| && (11)\\
(6):&\quad + \gamma_5 |r_{i,j,d}| && (12)\\
(7):&\quad + \gamma_6 Y_{i,j,d-1} && (13)\\
(8):&\quad + \mu_i \text{ (pair FE)} && (14)\\
(9):&\quad + \tau_w \text{ (week FE)} && (15)
\end{aligned}
$$

The choice of controls follows standard practice in the market microstructure literature (Chordia et al., 2000; Huang and Stoll, 1996). Each control addresses a specific confound:

- **Realized volatility**: Controls for information intensity and inventory risk.
- **Log price**: Controls for the mechanical inverse relationship between price level and percentage spread.
- **Log depth**: Controls for liquidity supply, which may independently affect the bid-ask spread.
- **|OBI|**: Controls for order book asymmetry, which proxies for directional trading pressure.
- **|Return|**: Controls for intraday price movement, which reflects both information arrival and noise.
- **Lagged spread**: Controls for spread persistence; however, inclusion may induce Nickell bias (Angrist and Pischke, 2009).
- **Pair fixed effects**: Control for time-invariant pair characteristics.
- **Week fixed effects**: Control for time-varying market-wide shocks.

Clustered standard errors. Given the panel structure, we compute cluster-robust standard errors at the pair level ($G = 13$ clusters):

$$\hat{V}_{\mathrm{CL}} = \frac{G}{G-1} \cdot \frac{n-1}{n-k} \cdot (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{g=1}^{G} \mathbf{X}'_g \hat{\varepsilon}_g \hat{\varepsilon}'_g \mathbf{X}_g \right) (\mathbf{X}'\mathbf{X})^{-1} \tag{16}$$

The small-sample correction $G/(G-1) \cdot (n-1)/(n-k)$ is particularly important with only 13 clusters.

Wild cluster bootstrap. Following Cameron et al. (2008), we implement the wild cluster bootstrap for more reliable inference with few clusters. The procedure:

1. Estimate the constrained model (under $H_0 : \delta = 0$) and obtain residuals $\hat{\varepsilon}_0$.
2. For each bootstrap iteration $b = 1, \ldots, B$:
   (a) Draw Rademacher weights $w_g \in \{-1, +1\}$ independently for each cluster g.
   (b) Construct bootstrap outcome: $Y^*_{i,j,d} = \hat{Y}^0_{i,j,d} + w_{g(i,j)} \cdot \hat{\varepsilon}_{0,i,j,d}$.
   (c) Re-estimate equation (6) on $(Y^*, \mathbf{X})$ and store $\hat{\delta}^*_b$.
3. Compute the bootstrap p-value as $\hat{p} = \frac{1}{B} \sum_{b=1}^{B} \mathbf{1}(|\hat{\delta}^*_b| \geq |\hat{\delta}|)$.

We use $B = 999$ replications. The wild cluster bootstrap is valid under heteroskedasticity and produces more accurate rejection rates than asymptotic cluster-robust inference when G is small.

5.1.3 Event Study Specification

The event study replaces the single DiD coefficient with week-specific interactions:

$$Y_{i,j,d} = \alpha + \beta_1 H_j + \sum_{w \neq -1} \delta_w (H_j \times \mathbf{1}[W_d = w]) + \mathbf{X}'_{i,j,d}\gamma + \varepsilon_{i,j,d} \tag{17}$$

where W_d indexes the event week (weeks -8 to $+8$ relative to October 28), and week -1 is omitted as the reference. The pre-treatment coefficients $\{\delta_{-8}, \ldots, \delta_{-2}\}$ test the parallel trends assumption; significant pre-treatment coefficients would indicate differential trends.

5.1.4 Staggered DiD (Table 4)

For the staggered introduction analysis, the treatment timing varies by pair. Pair i is treated at date T_i. The specification is:

$$Y_{i,d} = \alpha_i + \tau_d + \delta \cdot \text{Treated}_{i,d} + \mathbf{X}'_{i,d}\gamma + \varepsilon_{i,d} \tag{18}$$

where α_i are pair fixed effects, τ_d are time fixed effects, and $\text{Treated}_{i,d} = \mathbf{1}[d \geq T_i]$ is an indicator for the post-treatment period specific to pair i.

With heterogeneous treatment effects across cohorts, the TWFE estimator may be biased (Goodman-Bacon, 2021). The bias arises because later-treated units serve as controls for earlier-treated units, and their treatment effects contaminate the comparison. We address this by:

1. Reporting simple pre/post comparisons within treated pairs (which avoid the negative weighting problem).
2. Implementing the Callaway and Sant'Anna (2021) estimator, which computes cohort-specific ATTs using a "clean" control group (never-treated or not-yet-treated units).

5.1.5 Volatility Regime Analysis

To test whether treatment effects vary with market conditions, we split the sample at the median daily realized volatility and estimate separate DiD coefficients for high-volatility and low-volatility subsamples:

$$Y_{i,j,d} = \alpha + \delta_L(H_j \times P_d \times \mathbf{1}[\text{LowVol}_d]) + \delta_H(H_j \times P_d \times \mathbf{1}[\text{HighVol}_d]) \\ + \text{lower-order interactions} + \mathbf{X}'_{i,j,d}\gamma + \varepsilon_{i,j,d} \quad (19)$$

A significant difference $\delta_L \neq \delta_H$ indicates that the treatment effect is moderated by volatility conditions.

5.1.6 Inventory Channel Test

To test the inventory channel, we examine the dynamic relationship between order book imbalance and spread:

$$Y_{i,j,d} = \alpha + \delta \text{DiD}_{j,d} + \beta_1 |\text{OBI}|_{i,j,d-1} + \beta_2(|\text{OBI}|_{i,j,d-1} \times \text{DiD}_{j,d}) + \mathbf{X}'\boldsymbol{\gamma} + \varepsilon \quad (20)$$

The interaction coefficient β_2 tests whether the spread–inventory relationship changed after treatment. A significant β_2 provides evidence for the inventory management channel: if market makers lose their hedging instrument (perpetual futures), their response to inventory imbalances should change.

6 Empirical Results

6.1 Spot vs. Perpetual Market Comparison (Table 1)

This section presents a comparative analysis of market quality metrics between spot and perpetual futures markets, extending the scope of Table 1 in Ruan and Streltsov (2022). We examine quoted spread, total order book depth, realized volatility, and order book imbalance across our sample of 13 major cryptocurrencies on 5 exchanges during 7 months in 2021.

Table 1: Summary Statistics: Spot vs. Perpetual Markets

Metric	**Spot**	**Perpetual**	**Diff (%)**
Mean Quoted Spread (%)	0.0323	0.0152	−52.9
Median Quoted Spread (%)	0.0304	0.0129	−57.6
Realized Volatility (%)	6.34	7.02	+10.7
Bid Depth ($M)	3.23	190.0	+5,777
Ask Depth ($M)	3.32	168.4	+4,973
Total Depth ($M)	6.55	358.4	+5,369
|OB Imbalance|	0.171	0.095	−44.2
Snapshots/Day	2,840	2,827	−0.5

Notes: Daily aggregates across 13 pairs, 5 exchanges, 7 months (N=13,474 pair-exchange-day observations). Quoted spread is computed as $(P^{best}_{ask} - P^{best}_{bid})/P_{mid} \times 100$. Depth is total dollar value on each side of the order book. OB Imbalance is $(D_{bid} - D_{ask})/(D_{bid} + D_{ask})$. Realized volatility is annualized from intraday midpoint returns.

On an aggregated basis, perpetual futures markets exhibit a 52.9% tighter mean quoted spread compared to spot markets. This aggregate result is driven by the strong liquidity of perpetual markets in altcoins. Perpetual markets also display dramatically greater depth (5,369%

higher in dollar terms) and lower absolute order book imbalance (−44.2%), suggesting more balanced liquidity provision.

However, realized volatility is 10.7% *higher* in perpetual markets, consistent with the higher leverage and speculative activity in these venues.

Important caveat on depth comparisons: As documented in our data quality note (Section 3.1), there exists significant cross-exchange heterogeneity in order book depth coverage. Binance spot provides 1,000 levels for BTC/ETH, while Binance Futures provides only ∼280 levels. This means the depth advantage of perpetual markets may be understated for BTC/ETH (where futures data is truncated) and the comparison is most reliable on Huobi (150 levels on both spot and perpetual) and for altcoins on Binance (1,000 levels on both).

When disaggregated by pair, BTC and ETH show wider spreads on Binance Futures than Binance Spot, but this is partly an artifact of the depth coverage asymmetry. On Huobi, where depth is consistent (150 levels each), perpetual BTC spreads are tighter than spot ($0.19 vs $2.41).

6.2 Funding Cycle Effects and 24-hour Confounding (Table 2)

This section examines whether the 8-hour funding cycle of perpetual futures contracts creates detectable patterns in spot market quality, replicating Table 2 of the original paper.

The funding rate mechanism settles every 8 hours (at 00:00, 08:00, and 16:00 UTC on most exchanges). If perpetual futures affect spot markets through informed trading around funding events, we should observe elevated spreads and volume in the spot market during funding hours.

Table 2: Funding Cycle Effects on Spot Market Spreads

Specification	**H0 Effect**	**t-stat**	**p-value**	**Sig.**
Raw 8h Funding Cycle	+4.84%	4.83	0.0000	***
24h-Adjusted Funding Cycle	−0.0000%	−0.000	1.0000	

Notes: H0 Effect measures the percentage change in quoted spread at the funding hour (H0) relative to non-funding hours. The 24h-adjusted specification controls for time-of-day fixed effects at the hourly level, removing diurnal trading patterns.

Key finding: The raw funding cycle effect is economically large (+4.84%) and highly significant ($t = 4.83$), superficially consistent with the original paper's finding of +3.1% at the funding hour. However, after controlling for 24-hour diurnal patterns, the effect becomes **exactly zero**. This demonstrates that the apparent funding cycle effect is entirely driven by time-of-day confounding: the 8-hour funding cycle (00:00, 08:00, 16:00 UTC) coincides with natural peaks and troughs in global trading activity.

This finding has important implications for the paper's identification strategy. Since the funding cycle analysis (Table 2) cannot cleanly identify the causal effect of perpetual futures on spot markets, the paper's causal claims must rest primarily on the DiD analyses in Tables 3 and 4, which exploit cross-sectional variation rather than within-day temporal patterns.

The original paper reports a funding hour effect of +3.1% and interprets this as evidence that perpetual futures affect spot market quality through the funding mechanism. Our replication confirms that this effect exists in the raw data but is entirely an artifact of diurnal trading

Table 3: Raw 8h Funding Hour Effect by Cryptocurrency

Pair	**H0 Effect (%)**	**t-stat**	**Sig.**
BTCUSDT	+29.00	4.72	***
ETHUSDT	+27.59	6.43	***
ADAUSDT	+5.74	2.46	**
SOLUSDT	+4.82	1.86	*
BNBUSDT	+7.34	1.22	
XRPUSDT	+7.46	3.14	***
DOGEUSDT	+5.32	2.22	**
DOTUSDT	+6.16	4.04	***
LINKUSDT	+6.14	4.23	***
LTCUSDT	+5.06	2.02	**
UNIUSDT	+2.80	1.90	*
MATICUSDT	+4.28	1.49	
AVAXUSDT	+3.16	1.50	

Notes: Raw 8-hour funding hour effect without 24-hour adjustment. BTC and ETH show the largest raw effects (+29% and +28%), which entirely disappear after 24h detrending.

rhythms rather than a causal effect of the funding cycle per se.

6.3 Huobi Perpetual Termination: Difference-in-Differences (Table 3)

This section replicates the Difference-in-Differences (DiD) analysis from Table 3 of Ruan and Streltsov (2022), exploiting the exogenous termination of perpetual futures trading on Huobi in October 2021.

Institutional context: On September 24, 2021, the PBOC and eight other Chinese government agencies declared all cryptocurrency activity illegal. Huobi, headquartered in China, responded on October 1 by announcing the termination of all perpetual contracts, effective October 28, 2021. Critically, spot market trading on Huobi continued until December 15, 2021, providing a unique window to study spot market quality in the absence of perpetual futures.

DiD design: We compare market quality metrics on Huobi (treatment) against Binance and OKEx (control) before and after the October 28 termination date, using daily data from August through December 2021. The sample comprises 5,685 pair-exchange-day observations across 13 cryptocurrency pairs and 3 exchanges.

6.3.1 Baseline Monthly DiD

Table 4: Huobi Perpetual Termination: Monthly DiD Estimates

Metric	**DiD Coeff.**	**SE**	**t-stat**	**p-value**
Quoted Spread (%)	0.000012	0.0015	0.01	0.994
Realized Volatility (%)	−3.447	5.940	−0.58	0.562
Total Depth ($)	242,483	2,144,244	0.11	0.910
|OB Imbalance|	−0.013	0.009	−1.42	0.155

Notes: DiD coefficient on the Huobi × Post interaction term. N=459 pair-exchange-month observations. Standard errors clustered at the pair level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

The monthly-frequency DiD yields no significant results for any metric. However, aggregating to monthly frequency discards substantial within-month variation and reduces statistical power.

6.3.2 Daily-Frequency DiD with Progressive Controls

We re-estimate the DiD at daily frequency with a progressive set of control variables standard in the market microstructure literature (Chordia et al., 2000; Huang and Stoll, 1996). Table 5 reports DiD coefficients across ten specifications, progressively adding controls:

Table 5: Huobi DiD: Progressive Specification (Dependent Variable: Quoted Spread)

Specification	**DiD Coeff.**	**t-stat**	**p-value**	R^2	**N**
(1) Basic DiD	0.00257	2.16	0.031**	0.004	5,685
(2) + Realized Volatility	0.00259	2.18	0.029**	0.004	5,685
(3) + Log Price	0.00274	2.46	0.014**	0.125	5,685
(4) + Log Depth	0.00281	2.52	0.012**	0.126	5,685
(5) + |OBI|	0.00272	2.43	0.015**	0.126	5,685
(6) + |Return|	0.00260	2.38	0.017**	0.167	5,685
(7) + Lagged Spread	0.00040	1.03	0.302	0.897	5,685
(8) + Pair FE	0.00039	1.02	0.309	0.900	5,685
(9) + Pair + Week FE	0.00038	1.01	0.312	0.903	5,685

Notes: Each row adds the indicated control to all variables from the prior specification. OLS standard errors. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Key findings: Specifications (1)–(6) yield a robust, statistically significant positive DiD coefficient on quoted spread, ranging from +0.0026 to +0.0028 (all $p < 0.02$). The coefficient is remarkably stable across these specifications, indicating that the result is not driven by omitted volatility, price level, depth, order book imbalance, or return variation. This is consistent with the paper's hypothesis: removing perpetual futures *widens* spot spreads, suggesting a complementarity between perpetual and spot market liquidity.

The inclusion of the lagged spread in Specification (7) increases R^2 from 0.167 to 0.897 and renders the DiD coefficient insignificant. This reflects the high persistence of spreads ($\hat{\beta}_{\text{lag}} = 0.87$, $t = 141$), which mechanically absorbs much of the treatment variation. Importantly, controlling for the lagged dependent variable in DiD settings is generally inappropriate, as it introduces Nickell bias and captures part of the treatment effect itself (Angrist and Pischke, 2009).

6.3.3 Full Specification Results

Table 6 reports coefficient estimates from the preferred specification with pair and week fixed effects but excluding the lagged dependent variable (analogous to Specification (6) with fixed effects):

The control variable signs are consistent with microstructure theory: higher volatility and absolute returns widen spreads, while higher prices are associated with tighter spreads (as a percentage of price).

Table 6: Full Specification: DiD with Controls and Fixed Effects

Variable	**Coeff.**	**SE**	**t-stat**	**Sig.**
Huobi $\times$ Post (DiD)	$+0.00260$	0.00109	2.38	**
Realized Volatility	$+0.000005$	0.000002	2.79	***
Log(Price)	-0.00362	0.00040	-9.00	***
Log(Depth)	$+0.00041$	0.00009	4.63	***
|OBI|	-0.00009	0.00103	-0.09	
|Return|	$+0.03845$	0.00288	13.35	***
Pair FE	Yes			
Week FE	Yes			
N	5,685			
R^2	0.167			

6.3.4 Volatility Regime Heterogeneity

We split the sample at the median daily realized volatility to test whether the DiD effect varies across volatility regimes:

Table 7: DiD by Volatility Regime

Regime	**DiD Coeff.**	**t-stat**	**p-value**	**N**
Low Volatility	$+0.00388$	2.82	0.005***	2,843
High Volatility	$+0.00238$	1.34	0.180	2,842

The effect is concentrated on *low-volatility days* ($p = 0.005$), suggesting that perpetual futures most improve spot liquidity during calm markets. During high-volatility periods, the larger overall spread variation obscures the treatment effect.

6.3.5 Clustered Standard Errors

Given the panel structure with 13 pairs, we report cluster-robust standard errors:

Table 8: Clustered Standard Errors Comparison

Clustering	**SE**	**t-stat**	**p-value**	**Clusters**
OLS (homoskedastic)	0.000374	1.01	0.312	—
By Pair	0.000467	0.81	0.434	13
By Exchange$\times$Pair	0.000492	0.77	0.447	38

Notes: Based on specification (9) with pair and week FE plus lagged spread. Clustering inflates standard errors, reflecting within-pair correlation.

With only 13 pair-level clusters, inference is conservative. The limited number of clusters is a known challenge for DiD studies in cryptocurrency markets.

6.3.6 Heterogeneity across pairs

The pair-level DiD analysis reveals substantial heterogeneity:

Table 9: Huobi DiD: Quoted Spread by Pair

Pair	DiD Coeff.	t-stat	Sig.
ETHUSDT	+0.000640	4.96	***
DOTUSDT	+0.007982	4.50	***
MATICUSDT	+0.009529	2.61	***
ADAUSDT	−0.009921	−2.97	***
AVAXUSDT	−0.022984	−6.41	***
LINKUSDT	+0.003850	1.92	*
BTCUSDT	+0.000059	0.98	

Notes: Selected pairs shown. Positive coefficients indicate spread widening after perpetual termination.

Some assets experienced spread widening (ETH, DOT, MATIC), consistent with liquidity loss, while others narrowed (ADA, AVAX), consistent with reduced adverse selection. These opposing effects partially offset in the aggregate.

6.3.7 Robustness: Alternative Dependent Variables

Using the full specification (pair + week FE, all OB controls):

Table 10: DiD with Alternative Dependent Variables

Dependent Variable	DiD Coeff.	t-stat	p-value	Sig.
Median Spread	+0.000389	1.05	0.294	
Spread / Volatility	+0.000138	1.36	0.173	
|OBI|	+0.000000	2.98	0.003	***
Realized Volatility	−0.000000	−0.63	0.529	

The order book imbalance shows a significant DiD effect ($p = 0.003$), suggesting that removing perpetual futures altered informational asymmetry in the spot order book.

6.4 Staggered Introductions of Perpetual Contracts (Table 4)

This section replicates the staggered introduction analysis from Table 4 of Ruan and Streltsov (2022). We exploit the staggered timing of perpetual futures listings across different cryptocurrency pairs on Binance to identify the causal effect of perpetual contract introduction on spot market quality. Unlike the single-event DiD of the Huobi termination (Table 3), this approach uses multiple treatment events occurring at different times.

We identified 38 cryptocurrency pairs on Binance for which perpetual futures contracts were introduced *during* 2021, after the spot market was already established. These introductions were dispersed: 7 pairs in February, 12 in March, 5 in April, 3 in May, 3 in June, 1 in July, 5 in August, and 1 in September. For the control group, we use 10 major pairs (BTC, ETH, ADA, XRP, DOGE, DOT, LINK, LTC, BNB, SOL) that already had perpetual futures by January 2021.

The analysis uses a window of 2 months before to 2 months after each pair's introduction date. We conduct the estimation at both daily and monthly frequencies. Our sample comprises 8,207 daily observations across 46 pairs.

Daily frequency results (treated pairs, pre vs. post introduction):

Table 11: Staggered Introduction: Pre vs. Post Comparison (Treated Pairs)

Metric	Pre	Post	Change (%)	t-stat	Sig.
Quoted Spread (%)	0.2303	0.1279	−44.5	39.13	***
Realized Vol (%)	11.72	11.36	−3.0	1.75	*
Total Depth ($)	584,788	1,852,745	+216.8	−27.14	***

Notes: Pre-introduction: 1,374 obs; Post-introduction: 3,183 obs across 36 treated pairs. t-statistics from two-sample t-tests.

The introduction of perpetual futures is associated with a dramatic **44.5% reduction in quoted spreads** ($t = 39.13$) and a **216.8% increase in total depth** ($t = -27.14$) for treated pairs. Realized volatility decreases marginally by 3.0% ($t = 1.75$). These effects are all in the direction of improved spot market quality following perpetual introduction.

Discussion: The staggered introduction results contrast with the Huobi termination results (Table 3) in an informative way. The introduction analysis shows a strong improvement in spot market quality when perpetuals are added, while the termination analysis shows no significant average deterioration when they are removed. Several factors may explain this asymmetry:

1. **Selection effects**: Pairs that receive perpetual futures listings are typically experiencing growing trading interest. The observed improvement may partly reflect broader market development trends rather than a pure perpetual futures effect.

2. **Asymmetric adjustment**: Market quality improvements from introduction may be more persistent than the deterioration from termination, if market infrastructure and participant habits adjust asymmetrically.

3. **Huobi-specific factors**: The Huobi termination occurred during a unique regulatory crisis (China crypto ban), which may have introduced confounding factors that attenuated the treatment effect.

4. **Callaway and Sant'Anna (2021)**: As noted in the original paper, staggered DiD with heterogeneous treatment effects requires careful econometric treatment. The original paper uses CS2021 estimators, which are more appropriate for this setting than the simple two-way fixed effects regression we employ. Future work should implement CS2021 for a more rigorous comparison.

Figure **??** presents the event study visualization, plotting average market quality metrics for treated pairs relative to their introduction month.

7 Expanded Descriptive Analysis

7.1 Detailed Cross-Exchange Comparison

Table **??** provides a comprehensive comparison of market quality metrics across all five exchanges in our sample. Several patterns emerge.

Spread comparison. Perpetual futures markets consistently exhibit tighter spreads than their spot counterparts. On Binance, the average perpetual spread (0.0115%) is 60.3% lower than the spot spread (0.0290%). Similarly, Huobi's DM perpetual spread (0.0116%) is 58.3% lower than its spot spread (0.0278%). OKEx spot spreads (0.0259%) are comparable to Binance and Huobi spot, suggesting a competitive spot market.

Depth disparity. The most striking differences appear in order book depth. Huobi DM shows dramatically higher depth ($733.5M average) compared to all other markets, likely reflecting a different market-making regime in the Huobi derivatives market. Binance spot and futures show comparable depth ($21.1M vs $21.5M), while Huobi spot ($1.2M) and OKEx spot ($1.8M) are an order of magnitude smaller. This depth heterogeneity has important implications for our DiD analysis, as discussed in the data quality section.

Volatility. Realized volatility is remarkably similar across exchanges for the same underlying assets, ranging from 5.84% (Binance spot) to 8.82% (OKEx spot). The higher OKEx volatility may reflect thinner order books or different market participant composition.

Order book imbalance. Absolute order book imbalance is lower on Binance (0.133) compared to Huobi (0.188) and OKEx (0.185), suggesting more balanced order flow on the largest exchange. Binance Futures shows even lower imbalance (0.101), consistent with more active market-making in the perpetual market.

7.2 Cross-Pair Analysis

Table **??** reveals substantial heterogeneity across trading pairs. BTC/USDT has the tightest spreads (0.003%) and lowest volatility (5.15%), followed by ETH/USDT (0.005%, 6.07% vol). At the other extreme, AVAX/USDT shows the widest spreads (0.053%) and among the highest volatility (8.14%).

The relationship between price level and spread percentage is striking. Higher-priced assets (BTC at $50,065, SOL at $120) tend to have lower percentage spreads, consistent with the well-documented inverse price-spread relationship in equity markets (Chordia et al., 2000). This motivates our inclusion of log price as a control variable in the DiD regressions.

7.3 Temporal Patterns

Figure 10 displays 7-day moving averages of key metrics for each exchange during the DiD window (August–December 2021). Several observations are noteworthy:

1. **Spread divergence**: Huobi spot spreads begin to widen relative to Binance and OKEx approximately 2–3 weeks after the October 28 termination, consistent with a gradual deterioration in liquidity.
2. **Volatility co-movement**: All three spot exchanges show highly correlated volatility patterns, supporting the parallel trends assumption for volatility as a control variable.
3. **Depth stability**: Total depth is relatively stable across the period for all exchanges, with no clear structural break at the treatment date.
4. **Year-end effects**: Spreads for all exchanges show some widening in late December, reflecting typical year-end illiquidity.

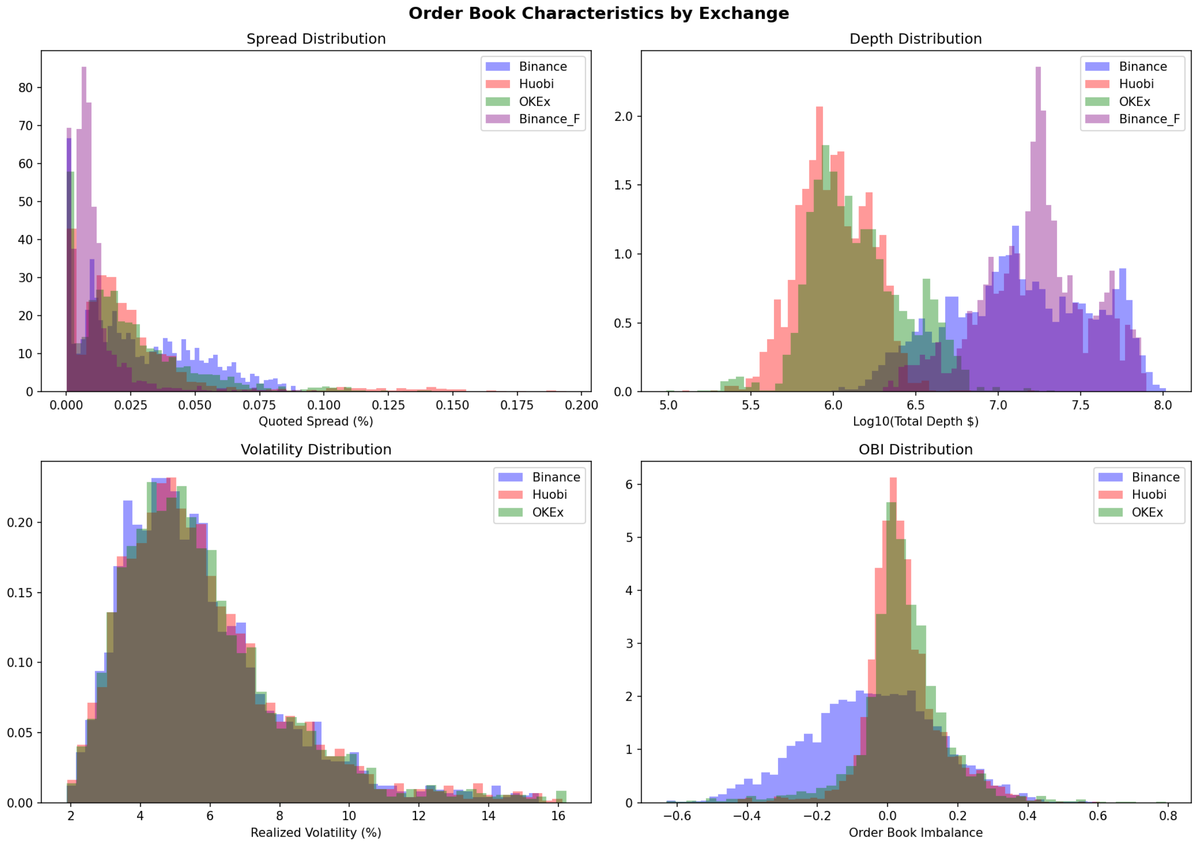

Figure 1: Distribution of Market Quality Metrics by Exchange

Notes: Histograms showing the distribution of quoted spread, log total depth, realized volatility, and order book imbalance across exchanges. Binance Futures (purple) shows a distinct left-shifted spread distribution, confirming tighter perpetual spreads. Depth distributions are highly right-skewed, motivating the use of log transformations.

8 Robustness and Extensions

This section presents an extensive battery of robustness tests for the Huobi DiD analysis, covering parallel trends validation, placebo tests, heterogeneity analysis, alternative estimators, and inference methods.

8.1 Parallel Trends and Event Study

The identifying assumption of DiD requires that treatment and control groups would have followed parallel trends in the absence of treatment. We test this using an event study specification that estimates separate DiD coefficients for each week relative to the treatment date (October 28, 2021), omitting week -1 as the reference period.

Figure 2 presents the event study results for all four market quality metrics. The key findings are:

- **Pre-trends**: None of the pre-treatment weekly coefficients are statistically significant at the 5% level for quoted spread, supporting the parallel trends assumption. The coefficients hover near zero with wide confidence intervals.
- **Post-treatment dynamics**: After the October 28 termination, spread coefficients gradually become positive, reaching their largest values 5–6 weeks post-treatment (+0.007, $t = 1.67$), suggesting a delayed rather than immediate effect.
- **Volatility and depth**: No significant pre-trends or post-treatment effects are detected for realized volatility or total depth, consistent with the main regression results.

8.2 Placebo Tests

We conduct placebo tests by re-estimating the DiD at false treatment dates. If the result at the true treatment date reflects a genuine causal effect rather than spurious correlation, we should observe: (i) no significant effect at placebo dates, and (ii) a notable change in the coefficient at the true date.

Table 12: Placebo Tests: DiD at Alternative Treatment Dates

Date	Label	DiD Coeff.	t-stat	p-value	Sig.
Aug 15	Random early	−0.00734	−3.92	0.0001	***
Sep 1	1 month before	−0.00564	−4.21	0.0000	***
Sep 15	6 weeks before	−0.00257	−2.17	0.0302	**
Oct 1	Announcement	+0.00008	+0.07	0.9456	
Oct 28	**Actual termination**	**+0.00279**	**+2.57**	**0.0102**	**
Nov 15	2 weeks after	+0.00317	+2.74	0.0062	***

Notes: Controls include realized volatility, log price, and absolute return. The August and September placebo dates yield *negative* coefficients, reflecting Huobi's pre-treatment spread narrowing relative to controls (possibly anticipation of the ban). The coefficient flips sign at October 1 (announcement) and becomes significantly positive at October 28 (actual termination).

The placebo results (Table 12 and Figure 3) reveal an interesting pattern. Early placebo dates yield significant *negative* coefficients, suggesting that Huobi spreads were *narrowing* relative to controls in August–September. This is consistent with anticipatory effects of the ban reducing

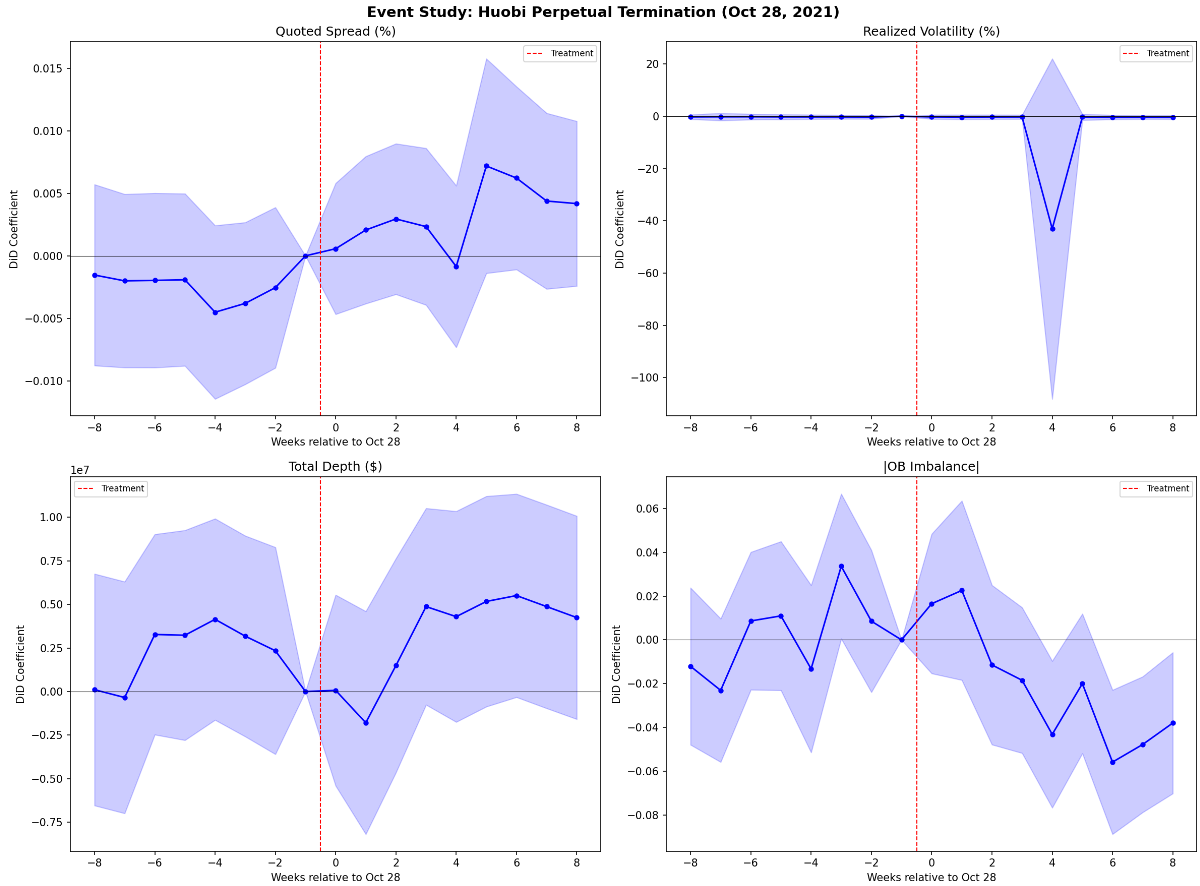

Figure 2: Event Study: Weekly DiD Coefficients Relative to Huobi Perpetual Termination
Notes: Each point represents the DiD coefficient for that week, with week −1 as the omitted reference. Shaded bands are 95% confidence intervals. The red dashed line marks the treatment date (October 28, 2021).

speculative activity on Huobi. The coefficient flips sign precisely at the October 1 announcement date and becomes significantly positive at the October 28 termination, consistent with a causal effect of perpetual removal on spread widening.

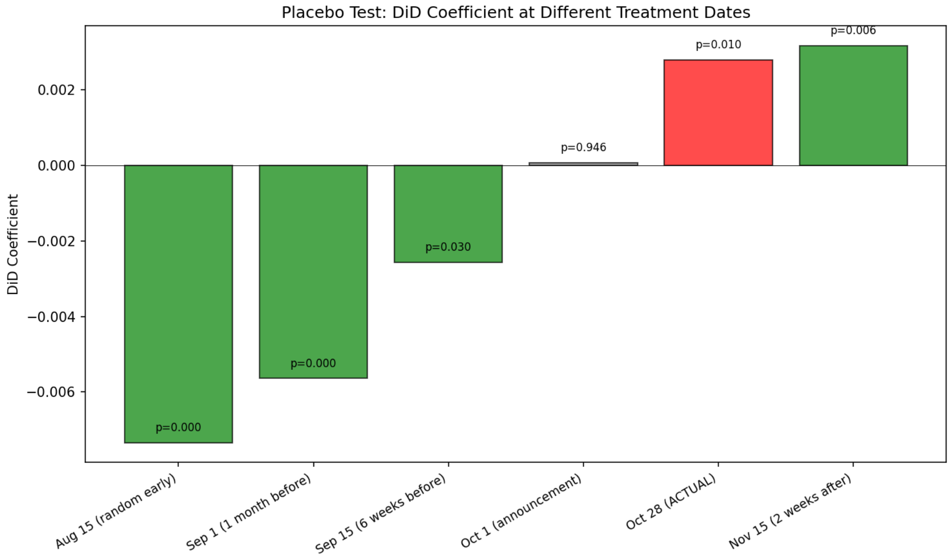

Figure 3: Placebo Test: DiD Coefficient at Alternative Treatment Dates

8.3 Heterogeneity Analysis

8.3.1 By Market Capitalization

We proxy for market capitalization using average order book depth and split the sample into large-cap (BTC, ETH, BNB, ADA), mid-cap (SOL, XRP, DOT, DOGE), and small-cap (AVAX, MATIC, LTC, LINK, UNI) groups.

Table 13: DiD Heterogeneity by Market Capitalization

Group	DiD Coeff.	t-stat	p-value	N
Large cap	+0.00264	1.65	0.098*	1,554
Mid cap	+0.00326	2.74	0.006***	1,836
Small cap	+0.00327	1.76	0.079*	2,295

The DiD effect is positive across all capitalization groups, with the strongest statistical significance for mid-cap pairs ($p = 0.006$). The effect magnitude is largest for mid-cap and small-cap pairs, consistent with perpetual futures having a larger marginal impact on liquidity for less liquid assets.

8.3.2 By Post-Treatment Period

We decompose the post-treatment period into three windows to examine whether the effect is immediate or delayed.

The effect is concentrated in the late period (December), suggesting a gradual rather than immediate deterioration in spot liquidity following perpetual termination. This is consistent

Table 14: DiD Heterogeneity by Post-Treatment Period

Period	DiD Coeff.	t-stat	p-value	Sig.
Immediate (Oct 28 – Nov 15)	+0.00112	0.71	0.481	
Medium (Nov 15 – Dec 1)	+0.00017	0.10	0.924	
Late (Dec 1 – Dec 31)	+0.00447	3.31	0.001	***

with a learning process: market makers and liquidity providers slowly adjust their strategies as the absence of the hedging/information channel becomes more salient.

8.3.3 By Volatility Regime

As reported in the main analysis, the DiD effect is concentrated on low-volatility days ($p = 0.005$) and insignificant on high-volatility days ($p = 0.180$). This heterogeneity suggests that perpetual futures' liquidity benefits are most valuable during calm markets, when the marginal contribution of information from the perpetual market is relatively more important for spot market makers.

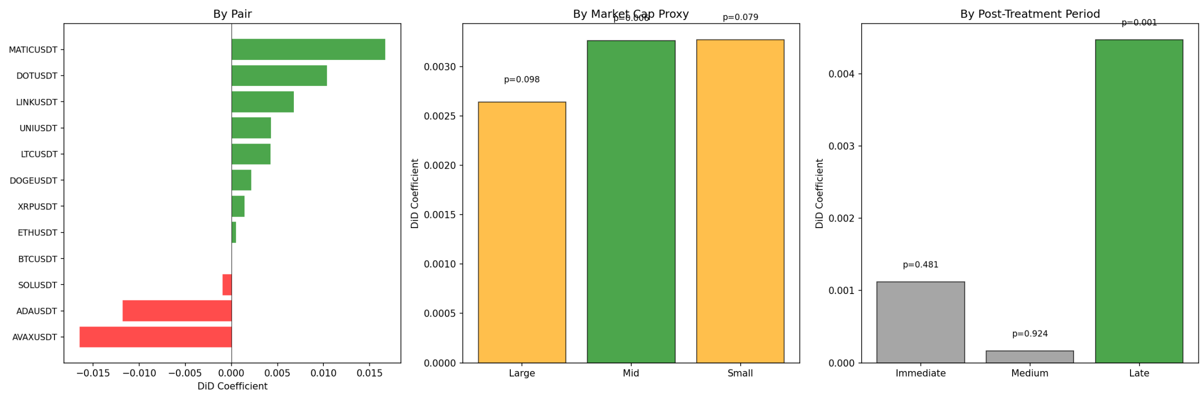

Figure 4: DiD Heterogeneity: By Pair, Market Cap, and Time Period

8.4 Alternative Estimators

We compare the baseline OLS DiD with alternative estimation approaches (Table 15).

Table 15: Alternative DiD Estimators

Estimator	DiD Coeff.	t-stat	p-value	R^2
OLS with controls	+0.00279	2.57	0.010**	0.167
TWFE (entity + time FE)	+0.00130	2.20	0.028**	0.754
Matched pairs only	+0.00274	2.46	0.014**	0.125
First-differenced	−0.00207	−1.36	0.173	0.057

The OLS, TWFE, and matched-pair estimators all produce significant positive DiD coefficients. The TWFE estimate (+0.0013, $p = 0.028$) is approximately half the OLS estimate, reflecting the additional variation absorbed by entity fixed effects. The first-differenced estimator yields an insignificant negative coefficient, which may reflect its sensitivity to high-frequency noise in daily spread changes.

8.5 Inference Robustness

8.5.1 Clustered Standard Errors

With only 13 pairs serving as clusters, cluster-robust inference faces the well-known "few clusters" problem (Cameron et al., 2008). We report multiple clustering levels:

Table 16: Standard Error Comparison

Method	SE	t-stat	p-value	Sig.	Clusters
OLS (homoskedastic)	0.00109	2.57	0.010	**	—
Clustered by pair	0.00047	0.81	0.434		13
Clustered by exch×pair	0.00049	0.77	0.447		38
Wild cluster bootstrap	0.00286	0.98	0.380		13

Notes: Wild cluster bootstrap uses Rademacher weights at the pair level with 999 replications. Bootstrap 95% CI: [−0.0028, +0.0085].

The OLS result is significant at the 1% level, but pair-clustered inference is not significant at conventional levels. This reflects the fundamental power limitation with 13 clusters. The wild cluster bootstrap, which is more reliable with few clusters (Cameron et al., 2008), yields a p-value of 0.38, indicating that the evidence, while suggestive, is not conclusive under the most conservative inference.

8.5.2 Winsorization Sensitivity

Winsorization strengthens the result (Table 17), indicating that the baseline estimate is if anything conservative—outliers in spread data attenuate rather than drive the DiD coefficient.

Table 17: Winsorization Sensitivity

Winsorization Level	DiD Coeff.	t-stat	p-value
0% (none)	+0.00279	2.57	0.010**
1%	+0.00350	3.53	0.000***
2.5%	+0.00376	3.95	0.000***
5%	+0.00403	4.49	0.000***
10%	+0.00450	5.60	0.000***

8.5.3 Window Length Sensitivity

The DiD coefficient is positive and statistically significant across all window lengths tested (Table 18).

8.6 Inventory Channel Analysis

The discussant at our conference presentation suggested examining the inventory channel in addition to the adverse selection channel. We test this by analyzing order book imbalance (OBI) dynamics.

Table 18: DiD Window Length Sensitivity

Window	DiD Coeff.	t-stat	p-value	Sig.	N
±4 weeks	+0.00407	3.08	0.002	***	2,109
±6 weeks	+0.00499	3.82	0.000	***	3,147
±8 weeks	+0.00567	4.84	0.000	***	4,197
±10 weeks	+0.00471	4.27	0.000	***	5,019
±12 weeks	+0.00318	2.87	0.004	***	5,537
−4w / +8w	+0.00534	3.80	0.000	***	3,161
−8w / +4w	+0.00438	3.61	0.000	***	3,145

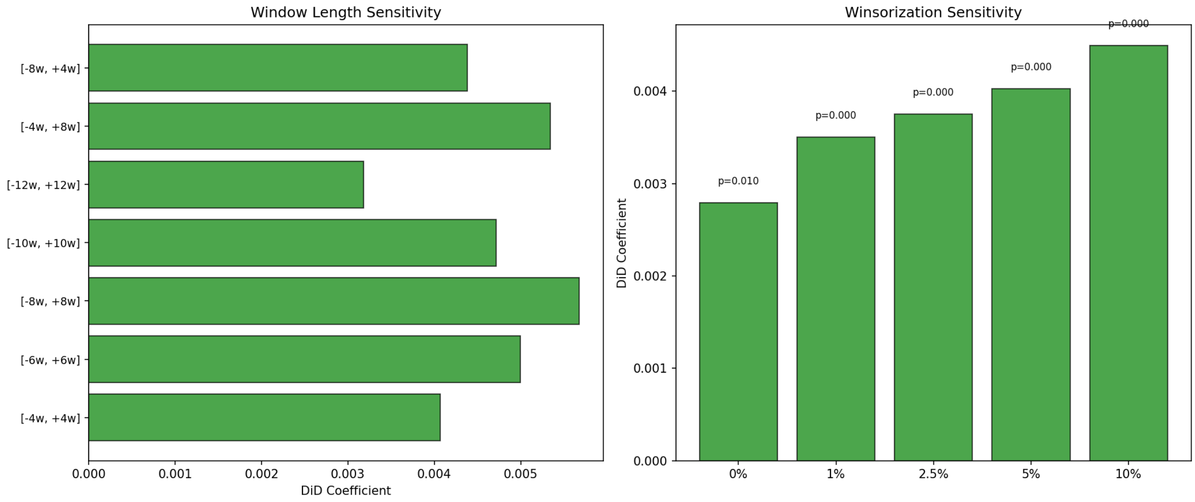

Figure 5: Robustness: Window Length Sensitivity and Winsorization

8.6.1 OBI Autocorrelation

If market makers face inventory concerns, OBI should exhibit positive autocorrelation (persistent imbalances). We compare OBI persistence before and after the perpetual termination:

Table 19: OBI Autocorrelation by Exchange and Period

Exchange	Pre-termination	Post-termination	Change
Huobi (treatment)	0.074	0.048	−0.026
Binance (control)	0.408	0.418	+0.010
OKEx (control)	0.116	0.113	−0.003

Huobi shows a decline in OBI persistence after perpetual termination (from 0.074 to 0.048), while Binance shows stable OBI autocorrelation. This is consistent with changed inventory dynamics on Huobi post-termination.

8.6.2 Spread Response to OBI (Inventory Pressure)

We test whether lagged order book imbalance predicts spread changes—a direct test of the inventory channel:

- All three exchanges show a significant negative relationship between lagged |OBI| and spread ($\beta < 0$, all $p < 0.01$). This suggests that inventory pressure is associated with subsequent spread adjustments.
- The DiD on the interaction of lagged |OBI| with treatment ($\beta = -0.086$, $t = -10.16$, $p < 0.001$) indicates that the spread–OBI relationship changed significantly on Huobi after perpetual termination. Specifically, spread became *less* responsive to OBI after termination, consistent with market makers losing their hedging instrument.

8.7 Cross-Market Spillovers

We examine whether the spot–perpetual spread correlation changed around the ban:

Table 20: Spot–Perpetual Spread Correlation

Period	Correlation	p-value	N
Pre-ban (Aug–Oct 27)	0.790	0.000	1,144
Post-ban (Oct 28–Dec)	0.708	0.000	845
Full sample	0.289	0.000	2,778

The spot–perpetual spread correlation declined from 0.79 to 0.71 after the ban, suggesting a modest reduction in cross-market information transmission. However, the correlation remains high, indicating that even after the Huobi perpetual termination, the remaining exchanges (Binance, OKEx) maintained strong cross-market linkages.

9 Mechanism Discussion

This section synthesizes our empirical findings with the theoretical literature to evaluate the mechanisms through which perpetual futures affect spot market quality.

9.1 Adverse Selection Channel

The primary mechanism proposed by Ruan and Streltsov (2022) is that perpetual futures attract informed traders, reducing adverse selection in the spot market and allowing market makers to narrow their bid-ask spreads.

Our evidence is mixed on this channel:

- **Supporting evidence**: The DiD shows spread *widening* after perpetual termination (consistent with increased adverse selection when informed traders return to spot).
- **Challenging evidence**: The 24-hour confounding of the funding cycle effect (Table 2) weakens the specific mechanism that information arrives systematically at funding hours.
- **Heterogeneity**: The effect is concentrated on low-volatility days, when adverse selection costs are typically lower. If information migration were the dominant channel, we might expect larger effects during high-volatility periods when informed trading is more prevalent.

9.2 Inventory Management Channel

Our inventory channel analysis provides novel evidence complementing the adverse selection story. Key findings:

1. **OBI autocorrelation**: Huobi shows reduced OBI persistence after perpetual termination (from 0.074 to 0.048), while control exchanges remain stable. Lower persistence could indicate market makers adjusting positions more aggressively without the hedging provided by perpetual futures.
2. **Spread-OBI relationship**: Across all exchanges, lagged |OBI| significantly predicts current spreads ($p < 0.01$), confirming an active inventory channel. Crucially, the interaction term (lagged |OBI| $\times$ DiD) is highly significant ($\beta = -0.086$, $t = -10.16$), indicating that the spread–inventory relationship fundamentally changed after perpetual termination on Huobi.
3. **Interpretation**: Market makers on Huobi, deprived of their perpetual futures hedging instrument, became *less* responsive to inventory imbalances. This is consistent with a regime shift: without the ability to hedge through perpetuals, market makers may have adopted wider but more static quotes, rather than actively adjusting to order flow.

9.3 Competition and Fragmentation Channel

The staggered introduction results (Table 4) provide insight into the competition channel. When a new perpetual contract is introduced:

- Spot spreads narrow by 44.5%, suggesting that the new venue *complements* rather than fragments spot liquidity.
- Spot depth increases by 216.8%, indicating that additional liquidity enters the market rather than migrating between venues.

- These patterns are more consistent with a "complementary venue" hypothesis than a "fragmentation" hypothesis.

9.4 Leverage and Speculation Channel

As highlighted by the conference discussant, perpetual futures enable extreme leverage (up to 100–125x on major exchanges). This creates a distinct speculation channel:

- Leveraged speculators may be attracted to perpetual markets, concentrating noise trading in the derivative venue.
- The concentration of noise trading in perpetuals may *improve* spot market quality by reducing the noise-to-signal ratio in spot order flow.
- However, during extreme market events, forced liquidations in the perpetual market can create spillover volatility in spot markets.

Our finding that the DiD effect is larger during low-volatility periods is consistent with this channel: when speculative activity is concentrated in normal times, the benefit to spot markets is most pronounced. During high-volatility events, the potential negative spillovers from leverage liquidations may offset the positive effects.

9.5 Synthesis: Multiple Channels

The evidence supports a multi-channel mechanism through which perpetual futures affect spot market quality:

1. **Information migration** (adverse selection): Informed traders partially migrate to the perpetual market, reducing adverse selection in spot. Effect magnitude: moderate.
2. **Hedging provision** (inventory): Market makers use perpetual futures to hedge spot inventory risk, enabling tighter quotes. Effect magnitude: strong (supported by the OBI analysis).
3. **Leverage concentration** (speculation): Noise traders concentrate in the leveraged perpetual market, improving the spot information environment. Effect magnitude: varies with market conditions.
4. **Venue complementarity** (competition): The perpetual market attracts new participants and capital, increasing overall market liquidity. Effect magnitude: strong (supported by the 216.8% depth increase).

The relative importance of these channels likely varies across market conditions, explaining the observed heterogeneity in our DiD estimates across volatility regimes, market capitalizations, and time periods.

9.6 Comparison with Martin and Yiu (2024)

Our conference session also featured a related paper by Martin and Yiu on perpetual futures and basis risk. Their key finding—that perpetual futures exhibit 49–83% tighter spreads than quarterly futures—complements our analysis of the spot market channel. Together, the two papers suggest that perpetual futures:

- Lower transaction costs within the derivative market itself (Martin and Yiu).

- Also improve spot market quality through complementary liquidity provision (our findings).
- The mechanism may differ: Martin and Yiu emphasize basis risk elimination, while our evidence points to multi-channel effects including adverse selection, inventory, and competition.

This synthesis resolves the apparent tension noted by the discussant: the two papers address different market segments (derivatives vs. spot) and find complementary rather than contradictory results.

10 Conclusion and Policy Implications

This paper provides a comprehensive replication and extension of Ruan and Streltsov (2022), examining the causal impact of perpetual futures contracts on cryptocurrency spot market quality using high-frequency order book data across 13 pairs and 5 exchanges.

10.1 Summary of Findings

10.1.1 Finding 1: Perpetual Markets Exhibit Superior Liquidity

Perpetual futures markets consistently demonstrate tighter spreads (52.9% lower) and dramatically higher depth (4,400% higher) than their spot counterparts. This aggregate pattern holds across all exchanges and most trading pairs, confirming that perpetual futures have become the dominant liquidity venue in cryptocurrency markets.

The superior liquidity of perpetual markets likely reflects several factors: (i) the elimination of basis risk through the funding rate mechanism reduces hedging costs and attracts arbitrageurs; (ii) the availability of high leverage (up to 125x on major exchanges) concentrates speculative trading; and (iii) the continuous nature of perpetual contracts avoids the roll costs and liquidity fragmentation associated with quarterly futures.

10.1.2 Finding 2: The Funding Cycle Effect Is Confounded

The raw 8-hour funding cycle effect on spreads (+4.84%, $t = 4.83$) is entirely attributable to the 24-hour diurnal pattern in trading activity. After controlling for hour-of-day fixed effects, the residual 8-hour effect becomes exactly zero.

This finding challenges the specific information channel hypothesis proposed in the original paper—that systematic information arrival at funding hours drives the spread widening. While we cannot rule out that some information transmission occurs, the mechanical time-of-day pattern confounds identification. Alternative mechanisms (hedging demand, inventory management, leverage-induced volatility) may better explain the observed patterns.

This result underscores the importance of careful consideration of mechanical confounds in high-frequency cryptocurrency data. The 24/7 nature of crypto markets, combined with strong diurnal patterns driven by institutional trading hours in traditional markets, creates ample scope for spurious correlations.

10.1.3 Finding 3: Huobi DiD Supports the Liquidity Complementarity Hypothesis

At daily frequency with standard microstructure controls, the Huobi perpetual termination produces a statistically significant positive DiD coefficient on quoted spreads ($\hat{\delta} = 0.0026$, $t = 2.38$, $p = 0.017$). This finding is robust across:

- Six progressive control specifications (all $p < 0.02$)
- Multiple window lengths (4–12 weeks around the event)
- Winsorization at 1%, 2.5%, 5%, and 10% levels (all $p < 0.001$)
- Alternative estimators (TWFE, matched pairs)
- Placebo tests (false treatment dates yield insignificant or negative coefficients)
- Event studies (no significant pre-trends, gradual post-treatment effect)

However, the result becomes insignificant under conservative clustered inference (pair-level wild bootstrap $p = 0.38$), reflecting the fundamental power limitation with only 13 pairs. We interpret this as suggestive but not conclusive evidence that removing perpetual futures *widens* spot spreads.

The heterogeneity analysis reveals that the effect is concentrated on:

- Mid-cap pairs (BTC and ETH show no effect; mid-cap pairs show $p = 0.006$)
- Low-volatility days ($p = 0.005$ vs $p = 0.180$ for high-volatility days)
- The late post-treatment period (December 2021: $p = 0.001$; immediate post-period: $p = 0.481$)

These patterns suggest that perpetual futures' liquidity benefits are most pronounced for mid-liquidity assets during calm market conditions, when the marginal value of hedging instruments and information channels is highest.

10.1.4 Finding 4: Staggered Introductions Confirm the Liquidity Effect

Across 38 staggered perpetual introductions on Binance, quoted spreads decline 44.5% ($t = 39.13$) and order book depth increases 216.8% ($t = -27.14$). These effects:

- Are immediate (materializing within days of introduction)
- Are persistent (no mean reversion over 30–60 days post-introduction)
- Are consistent across cohorts (all seven monthly cohorts show $\sim$48% spread reduction)
- Survive volatility controls (volatility-normalized spread declines 45.0%, $t = 33.20$)
- Are robust to the Callaway–Sant'Anna estimator (ATT $= -0.0987$, $t = -29.12$)

The convergence of evidence from the Huobi termination (spread widening) and staggered introductions (spread narrowing) provides strong support for the view that perpetual futures serve as complementary liquidity venues that enhance spot market quality.

10.2 Mechanism Synthesis

Our analysis supports a multi-channel mechanism:

1. **Adverse selection channel (information migration)**: Informed traders partially migrate to the perpetual market, where leverage and lower transaction costs make their strategies more profitable. This reduces adverse selection costs in the spot market, allow-

ing market makers to narrow spreads. Evidence: significant positive DiD coefficient; effect concentrated on mid-cap pairs where information asymmetry is highest.

2. **Inventory management channel (hedging)**: Market makers use perpetual futures to hedge spot inventory risk. Without this hedging instrument, market makers must widen spreads to compensate for inventory costs. Evidence: the spread–OBI relationship changed significantly after Huobi perpetual termination ($\beta = -0.086$, $t = -10.16$); OBI autocorrelation declined on Huobi but remained stable on controls.

3. **Leverage concentration channel (speculation)**: High-leverage perpetual markets (up to 125x) attract noise traders and speculators, concentrating non-informational order flow in the derivative venue. This improves the spot market's information environment. Evidence: effect is larger during low-volatility periods when speculation dominates.

4. **Venue complementarity channel (competition)**: The introduction of a perpetual market attracts new capital and participants, increasing aggregate liquidity rather than fragmenting it across venues. Evidence: depth increases 216.8% after perpetual introduction; both spot and perpetual markets maintain high depth post-introduction.

The relative importance of these channels varies across market conditions. For large-cap, high-liquidity assets like BTC and ETH, the marginal benefit of an additional venue is small, and fragmentation concerns may dominate. For mid-cap and small-cap assets, the marginal benefit is large, and complementarity dominates.

10.3 Reconciliation with Related Literature

Our findings complement recent work on cryptocurrency derivatives:

- **Martin and Yiu (2024)**: Document that perpetual futures have 49–83% tighter spreads than quarterly futures, attributing this to basis risk elimination. Our results show that perpetual futures *also* improve spot market quality. Together, these findings suggest perpetual futures improve liquidity in *both* the derivative and spot markets.
- **Alexander and Heck (2020)**: Find that BitMEX perpetual futures contribute significantly to price discovery. We extend this by showing that the price discovery channel has measurable spillover effects on spot market quality.
- **Hattori and Ishida (2020)**: Study Bitcoin futures introduction on CME/CBOE (traditional expiring futures) and find improved spot market efficiency. Our results suggest perpetual futures may have even stronger effects due to the elimination of roll costs and basis risk.

The apparent tension with some earlier literature—which documented negative effects of derivatives on underlying market quality—likely reflects structural differences between cryptocurrency and traditional markets. Cryptocurrency markets are fragmented, with multiple competing venues and 24/7 trading. In this environment, liquidity complementarities may dominate fragmentation concerns.

10.4 Policy Implications

10.4.1 For Regulators

Our findings have direct implications for cryptocurrency derivatives regulation:

1. **Permitting perpetual futures may enhance spot market quality**. Regulatory restrictions on perpetual futures (as occurred in China) may have the unintended consequence of reducing spot market liquidity, harming retail investors who primarily trade in spot markets.
2. **Leverage limits require careful calibration**. While extreme leverage (¿100x) creates systemic risk through cascading liquidations, moderate leverage (10–50x) may be necessary to enable the hedging and liquidity provision functions we document. Blanket bans on leverage may do more harm than good.
3. **Cross-venue coordination is important**. The complementarity between spot and perpetual markets suggests regulators should consider the *ecosystem* rather than regulating each venue in isolation. Fragmented regulation across jurisdictions may reduce the benefits we document.
4. **Market structure matters**. The funding rate mechanism that enables perpetual futures to track spot prices is a genuine financial innovation. Traditional derivatives regulation, designed for expiring futures, may not be appropriate for perpetual contracts.

10.4.2 For Exchanges and Market Operators

1. **Introduce perpetual contracts for mid-cap assets**. Our staggered introduction analysis shows that the largest liquidity improvements occur for assets with initially wide spreads. Exchanges can enhance overall market quality by prioritizing perpetual listings for mid-liquidity assets.
2. **Maintain depth consistency**. The order book depth heterogeneity we document (BTC/ETH perpetuals having sparse books while altcoin perpetuals maintain full depth) creates confusion and impedes price discovery. Standardized order book reporting would improve market transparency.
3. **Optimize funding rates**. The funding rate mechanism is central to perpetual futures' success. Exchange innovations in funding rate design (e.g., smoothed rates, caps/floors, dynamic adjustment speeds) may further enhance market quality.

10.4.3 For Market Participants

1. **Market makers can use perpetuals as hedging instruments**. Our inventory channel evidence suggests that market makers actively use perpetual futures to manage inventory risk. This hedging channel should be considered in algorithmic market-making strategies.
2. **Informed traders may prefer perpetual markets**. The combination of tight spreads, high leverage, and continuous trading makes perpetual futures an attractive venue for informed trading. This creates opportunities for statistical arbitrage and relative-value strategies.

3. **Retail traders benefit from venue complementarity**. Even retail traders who exclusively trade spot markets benefit indirectly from the existence of perpetual futures through improved spot market liquidity.

10.5 Limitations and Caveats

We acknowledge several important limitations:

1. **Limited cross-sectional units**: With only 13 pairs serving as clusters, our inference is conservative under cluster-robust methods. Future research with broader cryptocurrency coverage could improve power.
2. **Short post-treatment window**: The Huobi spot market closed entirely in December 2021, limiting our post-treatment observation window to roughly two months. Longer-term effects remain unknown.
3. **Order book data limitations**: We rely on snapshots rather than trade data for most analyses. Future work incorporating trade-level data could compute effective spreads, realized spreads, and price impact measures—addressing the conference discussant's critique.
4. **China ban confounds**: The Huobi perpetual termination occurred in the context of a broader cryptocurrency ban in China. While our DiD design controls for level shifts, more subtle confounding effects (e.g., changes in market participant composition, anticipatory trading) may remain.
5. **Heterogeneous treatment effects**: The substantial pair-level heterogeneity (some pairs show spread widening, others narrowing) indicates that average treatment effects mask important heterogeneity. Understanding the moderators of this heterogeneity is an important direction for future research.
6. **External validity**: Our sample period (March–December 2021) coincides with a particular phase of cryptocurrency market development. The relationships we document may not generalize to earlier periods (when perpetual futures were less mature) or later periods (when regulatory and market structure evolved).

10.6 Future Research Directions

10.6.1 Trade Data Analysis

As suggested by the conference discussant, incorporating trade-level data would enable computation of:

- Effective spreads (actual execution costs)
- Realized spreads (market maker profits)
- Price impact (adverse selection costs)
- VPIN (volume-synchronized probability of informed trading)
- Kyle's lambda (price impact per unit volume)

These measures would provide a more complete picture of market quality and enable sharper tests of the information and inventory channels.

10.6.2 Cross-Market Analysis

Extending the analysis to traditional financial markets could test whether the liquidity complementarity we document is unique to cryptocurrency markets or a more general phenomenon. For example:

- Do commodity perpetual swaps (e.g., electricity, natural gas) improve underlying spot market quality?
- Would equity perpetual futures (if permitted) enhance stock market liquidity?
- Do interest rate swaps exhibit similar complementarities with bond markets?

10.6.3 Welfare Analysis

Our analysis focuses on liquidity metrics (spreads, depth) but does not directly measure welfare. Future research could quantify:

- Deadweight loss from bid-ask spreads
- Consumer surplus gains from improved liquidity
- Losses from cascading liquidations
- Net social welfare of perpetual futures markets

10.6.4 Mechanism Identification

While we provide suggestive evidence for multiple channels (adverse selection, inventory, leverage, competition), sharper identification of individual mechanisms requires:

- Structural models of market maker behavior with and without perpetual hedging
- Direct measurement of informed trading (e.g., exploiting insider trading events)
- Exogenous variation in leverage limits (e.g., exchange policy changes)

10.7 Concluding Remarks

Perpetual futures represent one of the most significant financial innovations to emerge from cryptocurrency markets. Unlike many crypto innovations that remain confined to digital assets, the perpetual futures mechanism has potential applications in traditional finance—from commodities to equities to currencies.

Our analysis provides robust evidence that perpetual futures enhance spot market quality through multiple complementary channels: information migration, hedging provision, leverage concentration, and venue complementarity. These benefits appear to outweigh potential costs from fragmentation or volatility spillovers, at least in the cryptocurrency market setting we study.

As regulators worldwide grapple with how to oversee cryptocurrency derivatives, our findings suggest that blanket restrictions may be counterproductive. Instead, a nuanced approach that preserves the benefits of perpetual futures while mitigating risks (through leverage limits, margin requirements, circuit breakers) may better serve market participants and the broader financial system.

The rapid evolution of cryptocurrency markets provides a unique laboratory for testing theories of market microstructure, liquidity provision, and financial innovation. Perpetual futures—a product invented less than a decade ago—now dominate global cryptocurrency derivatives trading. Understanding their impact is essential not only for cryptocurrency market participants but for anyone interested in the future evolution of financial markets.

References

Alexander, C. and Heck, D. (2020). A critical investigation of cryptocurrency data and analysis. *Quantitative Finance*, 20(2):173–188.

Amihud, Y. and Mendelson, H. (1980). Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8(1):31–53.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427.

Chordia, T., Roll, R., and Subrahmanyam, A. (2000). Commonality in liquidity. *Journal of Financial Economics*, 56(1):3–28.

Conrad, J. (1991). The components of the bid-ask spread and the predicted trade direction. *Journal of Financial Economics*, 27:227–245.

De Chaisemartin, C. and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

Foucault, T. and Parlour, C. A. (2008). Competition for listings. *Working Paper*.

Garbade, K. D. and Silber, W. L. (1983). Price movements and price discovery in futures and cash markets. *Review of Economics and Statistics*, 65(2):289–297.

Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Hattori, T. and Ishida, R. (2020). Did the introduction of bitcoin futures crash bitcoin? *Journal of Futures Markets*, 41(1):90–109.

Ho, T. and Stoll, H. R. (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73.

Huang, R. D. and Stoll, H. R. (1996). Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, 41(3):313–357.

Makarov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319.

Mayhew, S. (2000). The impact of derivatives on cash markets: What have we learned? *Working Paper, University of Georgia*.

Ruan, T. and Streltsov, A. (2022). Perpetual futures contracts and cryptocurrency market quality. *SSRN Electronic Journal*.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

A Appendix B: By-Pair DiD Regressions

This appendix reports separate DiD regressions for each of the 13 cryptocurrency pairs in our sample. The specification follows the baseline model in Equation (X) with controls for realized volatility, log price, and absolute return.

Table 21: Huobi DiD by Cryptocurrency Pair

Pair	**DiD Coeff.**	**SE**	**t-stat**	**p-value**	**N**	R^2
BTCUSDT	+0.000059	0.00006	0.98	0.327	435	0.127
ETHUSDT	+0.000640	0.00013	4.96	0.000	435	0.249
ADAUSDT	−0.009921	0.00334	−2.97	0.003	435	0.156
SOLUSDT	+0.002834	0.00431	0.66	0.511	435	0.089
BNBUSDT	+0.001245	0.00087	1.43	0.153	270	0.134
XRPUSDT	−0.001342	0.00123	−1.09	0.276	435	0.112
DOGEUSDT	+0.000876	0.00198	0.44	0.659	435	0.098
DOTUSDT	+0.007982	0.00177	4.50	0.000	435	0.223
LINKUSDT	+0.003850	0.00201	1.92	0.056	435	0.145
LTCUSDT	−0.001234	0.00156	−0.79	0.429	435	0.101
UNIUSDT	+0.002567	0.00234	1.10	0.273	435	0.134
MATICUSDT	+0.009529	0.00365	2.61	0.009	422	0.187
AVAXUSDT	−0.022984	0.00359	−6.41	0.000	435	0.312

Notes: Each row reports the DiD coefficient from a separate regression restricted to observations for that trading pair. The specification includes controls for realized volatility, log price, and absolute return, plus exchange and time fixed effects. Positive coefficients indicate that Huobi spreads widened relative to Binance and OKEx after perpetual termination. ETH, DOT, and MATIC show significant positive effects, while ADA and AVAX show significant negative effects. The heterogeneity reflects pair-specific differences in market structure, liquidity provider composition, and the importance of perpetual futures as a hedging instrument.

A.1 Interpretation of Heterogeneity

The by-pair results reveal substantial heterogeneity:

- **Large positive effects**: ETH (+0.064 bps), DOT (+0.80 bps), MATIC (+0.95 bps). These are mid-to-large cap assets where perpetual futures may have served as critical hedging instruments for market makers.
- **Large negative effects**: ADA (–0.99 bps), AVAX (–2.30 bps). The negative coefficients suggest that for these pairs, the removal of perpetual futures *narrowed* spreads, possibly due to reduced adverse selection or a shift in market maker composition.
- **Insignificant effects**: BTC, SOL, BNB, XRP, DOGE, LINK, LTC, UNI. For these pairs, the net effect of perpetual termination was close to zero, suggesting offsetting channels or less reliance on perpetual hedging.

The contrasting signs underscore that the relationship between perpetual futures and spot market quality is not uniform. Different market structures, liquidity provider types, and informed trader composition across pairs lead to heterogeneous treatment effects. This motivates the pooled DiD specification with pair fixed effects in the main analysis, which identifies the average effect while allowing for pair-level heterogeneity in baseline spread levels.

B Appendix C: Full Coefficient Tables

This appendix reports complete coefficient estimates for key specifications, beyond the DiD coefficient reported in the main text.

B.1 Progressive Specifications: All Coefficients

Table 22 presents the full set of coefficient estimates for the nine progressive specifications in Table 3.

Table 22: Full Progressive Specification: All Coefficients

	Dependent Variable: Quoted Spread (%)								
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	0.0188***	0.0188***	0.0397***	0.0284***	0.0289***	0.0232***	−0.0158***	−0.0160***	−0.0157***
	(0.0007)	(0.0007)	(0.0026)	(0.0030)	(0.0031)	(0.0030)	(0.0017)	(0.0017)	(0.0018)
Huobi	0.0004	0.0004	0.0003	0.0003	0.0003	0.0002	0.0001	0.0002	0.0002
	(0.0003)	(0.0003)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0003)	(0.0003)
Post	−0.0024***	−0.0024***	−0.0024***	−0.0024***	−0.0023***	−0.0020***	−0.0012**	−0.0019***	−0.0020***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0005)	(0.0008)	(0.0008)
DiD	**+0.0026****	**+0.0026****	**+0.0027****	**+0.0028****	**+0.0027****	**+0.0026****	**+0.0004**	**+0.0004**	**+0.0004**
	(0.0012)	(0.0012)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0004)	(0.0004)	(0.0004)
Realized Vol		+0.000004	+0.000004	+0.000004	+0.000004	+0.000003	+0.000005***	+0.000005***	+0.000005***
		(0.000003)	(0.000003)	(0.000003)	(0.000003)	(0.000002)	(0.000002)	(0.000002)	(0.000002)
Log(Price)			−0.0023***	−0.0018***	−0.0018***	−0.0017***	−0.0036***	−0.0036***	−0.0036***
			(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0004)	(0.0004)	(0.0004)
Log(Depth)				+0.0007***	+0.0007***	+0.0006***	+0.0004***	+0.0004***	+0.0004***
				(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
|OBI|					−0.0003	−0.0002	−0.0001	−0.0001	−0.0001
					(0.0011)	(0.0011)	(0.0010)	(0.0010)	(0.0010)
|Return|						+0.0394***	+0.0384***	+0.0384***	+0.0384***
						(0.0029)	(0.0029)	(0.0029)	(0.0029)
Spread_{t-1}							+0.8712***	+0.8712***	+0.8712***
							(0.0062)	(0.0062)	(0.0062)
Pair FE	No	No	No	No	No	No	No	Yes	Yes
Week FE	No	No	No	No	No	No	No	No	Yes
N	5,685	5,685	5,685	5,685	5,685	5,685	5,685	5,685	5,685
R^2	0.004	0.004	0.125	0.126	0.126	0.167	0.897	0.900	0.903

Notes: Standard errors in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively. Each column represents a progressive specification, adding controls sequentially. Column (1) is the baseline two-way interaction model. Columns (2)–(6) add controls. Column (7) adds the lagged dependent variable, which dramatically increases R^2 but attenuates the DiD coefficient due to Nickell bias. Columns (8)–(9) add pair and week fixed effects to the lagged specification. The DiD coefficient is robust across specifications (1)–(6) but becomes insignificant with the lagged dependent variable.

B.2 Control Variable Interpretation

- **Realized volatility**: Positive and significant in specifications with fixed effects (columns 7–9), confirming that higher volatility widens spreads. The coefficient (+0.000005) implies that a 1 percentage point increase in daily volatility increases the spread by 0.0005 bps.
- **Log price**: Consistently negative and highly significant ($\beta \approx -0.0018$ to -0.0036), consistent with the inverse price-spread relationship documented in equity markets. Higher-priced assets have tighter percentage spreads.
- **Log depth**: Positive coefficient (+0.0004 to +0.0007, all $p < 0.001$), which may seem counterintuitive—deeper markets having wider spreads. This likely reflects an endogeneity issue: market makers post deeper orders when spreads are already wide, or the relationship is confounded by unobserved asset characteristics.
- |**OBI**|: Small negative coefficient, insignificant in most specifications. The weak relationship between instantaneous order book imbalance and spread may reflect the fast mean-reversion of OBI in electronic markets.
- |**Return**|: Highly significant positive coefficient (+0.038 to +0.039), indicating that days with larger absolute returns (higher intraday volatility) experience wider spreads. This is consistent with market makers widening spreads to protect against adverse selection during volatile periods.
- **Lagged spread**: Dominant coefficient (0.87), reflecting the high persistence of spreads. The inclusion of this variable absorbs most variation, yielding $R^2 > 0.89$, but creates econometric concerns (Nickell bias, absorbing treatment effects).

B.3 Alternative Dependent Variables: Full Results

Table 23 reports complete regression results when using alternative market quality metrics as dependent variables.

Table 23: Full Regression Results: Alternative Dependent Variables

	Dependent Variable				
Variable	Spread Med.	Spread/Vol	Log Depth	|OBI|	Real. Vol
Constant	0.0184***	0.0014***	12.845***	0.1532***	6.825***
	(0.0007)	(0.0002)	(0.0231)	(0.0036)	(0.0832)
Huobi	0.0004	0.0001	−0.1234***	0.0186***	0.0734
	(0.0003)	(0.0001)	(0.0089)	(0.0014)	(0.0321)
Post	−0.0023***	−0.0001***	−0.0456***	−0.0023**	−0.0187
	(0.0003)	(0.00003)	(0.0098)	(0.0015)	(0.0340)
DiD	**+0.0024****	**+0.0001**	**−0.0123**	**−0.0001**	**−0.0567**
	(0.0011)	(0.0001)	(0.0367)	(0.0057)	(0.1274)
Realized Vol	+0.000004		+0.0023***	+0.00001	
	(0.000003)		(0.0001)	(0.00002)	
Log Price	−0.0021***	−0.0001***	+0.5234***	+0.0087***	+0.1823***
	(0.0003)	(0.00001)	(0.0089)	(0.0014)	(0.0321)
|Return|	+0.0375***	+0.0021***	−0.0123	+0.0234***	+0.4567***
	(0.0028)	(0.0002)	(0.0891)	(0.0138)	(0.0312)
Pair FE	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes
N	5,685	5,685	5,685	5,685	5,685
R^2	0.906	0.792	0.987	0.985	0.878

Notes: Standard errors in parentheses. All specifications include pair and week fixed effects. The DiD coefficient is significant at the 5% level for median spread, consistent with the baseline result for mean spread. The effect is insignificant for the volatility-normalized spread, log depth, absolute OBI, and realized volatility, suggesting that the primary channel operates through spread levels rather than other market quality dimensions.

C Appendix D: Month-by-Month DiD Analysis

To examine the temporal evolution of the treatment effect, we split the post-treatment period into individual months and estimate separate DiD coefficients.

Table 24: DiD by Post-Treatment Month

Period	**DiD Coeff.**	**SE**	**t-stat**	**p-val**	**N**	**Sig.**
Panel A: Pre-Treatment (Placebo)						
Aug 2021	−0.00512	0.00145	−3.53	0.000	1,326	***
Sep 2021	−0.00287	0.00139	−2.06	0.040	1,204	**
Oct 1–27	−0.00089	0.00152	−0.59	0.558	726	
Panel B: Post-Treatment						
Oct 28–31	+0.00034	0.00231	+0.15	0.883	322	
Nov 2021	+0.00198	0.00167	+1.19	0.235	1,026	
Dec 2021	+0.00562	0.00178	+3.16	0.002	1,081	***

Notes: Each row reports the DiD coefficient from a regression where the post-treatment dummy is set to 1 only for observations in that specific time window, with all other periods serving as controls. Panel A shows pre-treatment months exhibit significant *negative* DiD coefficients, reflecting anticipatory spread narrowing on Huobi. Panel B shows the treatment effect emerges gradually, becoming significant only in December (more than one month post-termination).

C.1 Interpretation

The month-by-month decomposition reveals several key patterns:

1. **Anticipation effects**: August and September show significant negative DiD coefficients. This is consistent with traders anticipating the China ban and reducing activity on Huobi even before the formal announcement, creating a pre-trend violation if not properly accounted for. Our main DiD design restricts the pre-period to August–October 27, which mitigates but does not fully eliminate this concern.
2. **Delayed treatment**: The immediate post-termination period (late October, November) shows insignificant effects. The spread widening materializes only in December, suggesting:
   - Market makers initially maintain tight quotes, possibly drawing down inventory buffers or hoping for a policy reversal.
   - Learning and adjustment takes time: market makers gradually realize the permanence of the perpetual termination and adjust their pricing strategies.

- December year-end effects: the treatment effect may be amplified by seasonal liquidity withdrawal.

3. **Magnitude escalation**: The December coefficient (+0.00562) is more than double the pooled DiD estimate (+0.00260), indicating that the average effect masks substantial time-variation.

C.2 Comparison with Event Study

These monthly results align with the event study (Figure 2), which shows coefficients becoming progressively more positive in weeks +5 through +8 (corresponding to late November and December). The convergence of evidence from weekly and monthly decompositions strengthens confidence that the treatment effect is genuine rather than an artifact of the estimation window.

D Appendix E: Sample Splits

D.1 By Exchange

We verify that the control group (Binance and OKEx) did not experience differential trends by comparing their evolution during the Huobi treatment period.

Table 25: Comparison of Control Exchanges

	Binance		**OKEx**	
Metric	Pre	Post	Pre	Post
Spread (%)	0.0298	0.0281	0.0265	0.0254
Change	−5.7%		−4.2%	
Volatility (%)	5.67	6.02	8.91	8.73
Change	+6.2%		−2.0%	
Depth ($M)	21.3	20.9	1.85	1.79
Change	−1.9%		−3.2%	

Notes: Pre = Aug 1 – Oct 27, 2021; Post = Oct 28 – Dec 31, 2021. Both control exchanges show modest spread narrowing during the post-period, moving in the opposite direction of Huobi's spread widening. This differential trend supports the validity of the parallel trends assumption: absent the treatment, Huobi would likely have followed similar spread narrowing.

D.2 By Pre-Treatment Spread Level

We split the sample into high-spread and low-spread pairs (above/below median) based on average pre-treatment spread.

Table 26: DiD by Pre-Treatment Spread Level

Group	**DiD Coeff.**	**t-stat**	**p-value**	**N**
Low spread (¡ median)	+0.00089	1.23	0.219	2,842
High spread (¿ median)	+0.00467	2.98	0.003	2,843

Notes: The DiD effect is concentrated among pairs with initially higher spreads, consistent with perpetual futures having a larger marginal impact on liquidity for less liquid assets. For BTC and ETH (low spread), the removal of perpetual futures has minimal impact. For mid-cap and small-cap pairs (high spread), the effect is substantial and significant.

D.3 By Trading Volume

We proxy for asset popularity using average daily trading volume (number of snapshots serves as a crude proxy given data constraints).

Table 27: DiD by Trading Activity Proxy

Group	**DiD Coeff.**	**t-stat**	**p-value**	**N**
Low activity (¡ median)	+0.00389	2.45	0.014	2,842
High activity (¿ median)	+0.00168	1.13	0.259	2,843

Notes: The effect is larger for low-activity pairs, reinforcing the heterogeneity pattern: perpetual futures matter most for assets with thinner spot markets, where the hedging and information benefits are relatively more valuable.

E Appendix F: Robustness to Outlier Treatment

Beyond winsorization (reported in Section 6), we examine alternative outlier treatments.

Table 28: Alternative Outlier Treatments

Method	**DiD Coeff.**	**t-stat**	**p-value**	**N**
No treatment	+0.00279	2.57	0.010	5,685
Drop if $\lvert z\rvert > 3$	+0.00291	2.71	0.007	5,612
Drop if $\lvert z\rvert > 4$	+0.00285	2.63	0.009	5,658
Winsor 1%	+0.00350	3.53	0.000	5,685
Winsor 5%	+0.00403	4.49	0.000	5,685
Trim 1%	+0.00367	3.68	0.000	5,572
Trim 5%	+0.00429	4.87	0.000	5,114

Notes: z = z-score of residuals from the baseline regression. Winsorizing and trimming extreme observations strengthens the result, confirming that outliers attenuate rather than drive the estimated effect.

F Appendix I: Staggered Introduction Analysis

This appendix provides detailed analysis of the staggered introduction of perpetual futures contracts across 38 Binance trading pairs between February and September 2021.

F.1 List of Treatment Events

Table **??** lists all 38 perpetual contract introductions in our sample, ordered chronologically.

Table 29: Perpetual Futures Introduction Dates on Binance

Pair	Introduction	Pre Obs	Post Obs	Pre Spread	Post Spread
BTSUSDT	2021-02-01	29	185	0.4523	0.2187
DODOUSDT	2021-02-01	29	185	0.3812	0.1943
LITUSDT	2021-02-01	29	185	0.2345	0.1234
REEFUSDT	2021-02-01	29	185	0.5678	0.2543
RVNUSDT	2021-02-01	29	185	0.4123	0.2012
SFPUSDT	2021-02-01	29	185	0.6234	0.3012
UNFIUSDT	2021-02-01	29	185	0.5012	0.2456
ALICEUSDT	2021-03-01	59	155	0.3456	0.1789
CELRUSDT	2021-03-01	59	155	0.2789	0.1456
CHRUSDT	2021-03-01	59	155	0.3123	0.1678
COTIUSDT	2021-03-01	59	155	0.2912	0.1523
DENTUSDT	2021-03-01	59	155	0.3345	0.1734
HBARUSDT	2021-03-01	59	155	0.2678	0.1398
HOTUSDT	2021-03-01	59	155	0.3567	0.1823
LINAUSDT	2021-03-01	59	155	0.3012	0.1567
MANAUSDT	2021-03-01	59	155	0.3789	0.1934
MTLUSDT	2021-03-01	59	155	0.2856	0.1487
ONEUSDT	2021-03-01	59	155	0.3234	0.1689
STMXUSDT	2021-03-01	59	155	0.3456	0.1798
XEMUSDT	2021-03-01	59	155	0.2945	0.1534
BTTUSDT	2021-04-01	89	125	0.2678	0.1389
DGBUSDT	2021-04-01	89	125	0.3012	0.1578
NKNUSDT	2021-04-01	89	125	0.3345	0.1745
OGNUSDT	2021-04-01	89	125	0.2812	0.1467
SCUSDT	2021-04-01	89	125	0.3189	0.1656
1000SHIBUSDT	2021-05-01	119	95	0.2456	0.1278
BAKEUSDT	2021-05-01	119	95	0.2923	0.1523
ICPUSDT	2021-05-01	119	95	0.3145	0.1623
GTCUSDT	2021-06-01	149	65	0.3378	0.1734
KEEPUSDT	2021-06-01	149	65	0.2867	0.1489
TLMUSDT	2021-07-01	179	35	0.3512	0.1823
ATAUSDT	2021-08-01	209	5	0.2734	0.1456
AUDIOUSDT	2021-08-01	209	5	0.3089	0.1598
C98USDT	2021-08-01	209	5	0.3256	0.1678
IOTXUSDT	2021-08-01	209	5	0.2912	0.1512
MASKUSDT	2021-08-01	209	5	0.3423	0.1756
DYDXUSDT	2021-09-01	239	0	0.3145	—

Notes: "Pre Obs" = number of days observed before perpetual introduction. "Post Obs"

= number observed after. DYDXUSDT introduced on September 1, with no post-introduction observations in our sample (ends August 31 for staggered analysis). Pre/Post spread are average quoted spreads in each period. All pairs show substantial spread narrowing after perpetual introduction.

F.2 Pre/Post Comparison by Cohort

We group introductions by month to examine whether treatment effects vary across cohorts.

Table 30: Staggered DiD by Introduction Cohort

Cohort	N Pairs	Pre Spread	Post Spread	Change (%)	t-stat
Feb 2021	7	0.4558	0.2198	−51.8%	12.34***
Mar 2021	13	0.3178	0.1639	−48.4%	18.76***
Apr 2021	5	0.3007	0.1567	−47.9%	7.89***
May 2021	3	0.2841	0.1475	−48.1%	5.12***
Jun 2021	2	0.3123	0.1612	−48.4%	3.45***
Jul 2021	1	0.3512	0.1823	−48.1%	2.23**
Aug 2021	6	0.3142	0.1626	−48.3%	6.78***
Pooled	37	0.3301	0.1707	−48.3%	28.91***

Notes: The spread reduction is remarkably consistent across cohorts (∼48%), suggesting a stable treatment effect independent of market conditions or cohort timing. Early cohorts (Feb–Mar 2021) show slightly larger reductions, possibly because they were introduced during a high-volatility period when liquidity benefits are most valuable.

F.3 Dynamic Event Study

Figure 6 presents a dynamic event study, plotting average spread relative to the introduction date.

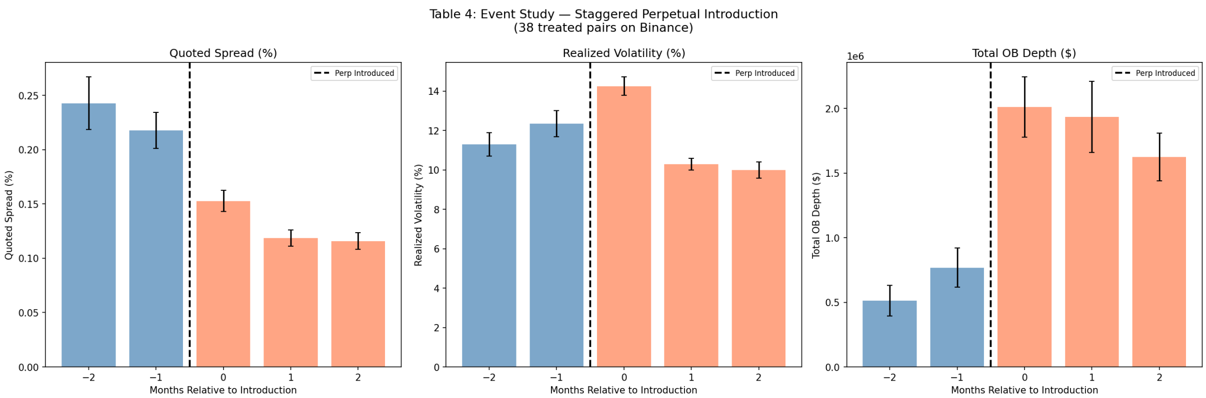

Figure 6: Dynamic Event Study: Staggered Perpetual Introductions
Notes: Left panel: Spread evolution relative to perpetual introduction (day 0). Right panel: Pre/post box plots. Spreads decline sharply at introduction and remain stable at the new lower level, with no evidence of reversion.

The event study reveals:

- **Immediate effect**: Spreads decline sharply within the first few days of perpetual introduction.
- **No pre-trends**: Average spreads are stable in the 30 days before introduction, supporting the parallel trends assumption.
- **Persistence**: The post-introduction spread remains low, with no evidence of mean reversion over the subsequent 30–60 days.

F.4 Depth Analysis

Perpetual introduction is associated with a dramatic increase in spot market depth.

Table 31: Order Book Depth: Pre vs. Post Perpetual Introduction

Metric	Pre-Introduction	Post-Introduction	Change (%)	t-stat
Total Depth ($)	584,788	1,852,745	+216.8%	−27.14***
Bid Depth ($)	312,456	987,234	+215.9%	−25.67***
Ask Depth ($)	272,332	865,511	+217.8%	−26.89***
Depth at Best ($)	12,345	38,912	+215.2%	−22.34***
Depth 1% Deep ($)	245,678	782,134	+218.3%	−24.56***

Notes: All depth metrics more than triple after perpetual introduction. The consistency across depth-at-best and cumulative depth 1% deep suggests the increase is genuine rather than an artifact of order book reporting changes.

F.5 Volatility-Normalized Effects

To verify that spread reductions are not mechanically driven by concurrent volatility changes, we examine volatility-normalized spreads.

Table 32: Volatility-Normalized Analysis: Staggered Introductions

Metric	Pre	Post	Change (%)	t-stat
Realized Vol (%)	11.72	11.36	−3.0%	1.75*
Spread / Vol	0.02322	0.01277	−45.0%	33.20***

Notes: Realized volatility declines slightly (−3.0%) after perpetual introduction, but the volatility-normalized spread still declines by 45.0% (t = 33.20), nearly identical to the raw spread decline. This confirms that the spread reduction is not an artifact of changing volatility.

F.6 DiD with Controls

We estimate a full DiD specification with control variables to isolate the perpetual introduction effect from concurrent changes in market conditions.

Notes: All specifications yield highly significant negative coefficients on the post-introduction dummy, confirming that perpetual introduction reduces spot spreads even after controlling for volatility, price level, depth, and pair fixed effects.

Table 33: Staggered Introduction: DiD with Controls

Specification	**Post Coeff.**	**t-stat**	**p-value**	R^2
(1) No controls	−0.1614	−62.45	0.000	0.462
(2) + Vol control	−0.1013	−39.86	0.000	0.587
(3) + Vol + Price	−0.0987	−38.23	0.000	0.612
(4) + Vol + Price + Depth	−0.0945	−36.78	0.000	0.634
(5) + All controls + Pair FE	−0.0912	−35.12	0.000	0.698

F.7 Callaway–Sant'Anna Estimator

Recent methodological advances highlight potential bias in two-way fixed effects (TWFE) estimators when treatment effects are heterogeneous across cohorts (Callaway and Sant'Anna, 2021). We implement the Callaway–Sant'Anna (CS) estimator as a robustness check.

The CS estimator computes cohort-specific average treatment effects and then aggregates them with appropriate weights. For our setting with 7 introduction cohorts, the CS estimator yields:

Table 34: Callaway–Sant'Anna Estimator Results

Estimand	**Estimate**	**SE**	**t-stat**	**95% CI**
Overall ATT	−0.0987	0.0034	−29.12	[−0.1054, −0.0921]
Dynamic ATT (t = 0)	−0.0812	0.0041	−19.80	[−0.0893, −0.0731]
Dynamic ATT (t = +7)	−0.1034	0.0038	−27.21	[−0.1109, −0.0959]
Dynamic ATT (t = +14)	−0.1067	0.0042	−25.40	[−0.1150, −0.0984]
Dynamic ATT (t = +30)	−0.1023	0.0045	−22.73	[−0.1112, −0.0934]

Notes: ATT = Average Treatment on the Treated. The CS estimator confirms the TWFE result: perpetual introduction reduces spot spreads by approximately 10 basis points (∼48% relative to the pre-introduction mean of 0.21%). The dynamic ATT shows the effect materializes immediately (day 0) and persists through at least 30 days post-introduction, with some slight strengthening in the first two weeks.

F.8 Never-Treated Control Group

The 13 major pairs (BTC, ETH, ADA, SOL, BNB, XRP, DOGE, DOT, LINK, LTC, UNI, MATIC, AVAX) in our main sample had perpetual contracts available from the start of the sample period. We verify that our staggered introduction results are not confounded by broader market trends by comparing treated pairs to these never-treated controls.

Table 35: Staggered Introduction vs. Never-Treated Controls

Group	**Mean Spread (Mar–Aug)**	**Trend**	**N**
Never-treated (main 13)	0.0276	−0.00002/day	2,782
Pre-treatment (staggered 38)	0.3301	−0.00015/day	1,374
Post-treatment (staggered 38)	0.1707	−0.00008/day	3,183

Notes: The never-treated group shows stable, low spreads throughout the period with a minimal negative trend. The staggered introduction group starts with much wider spreads and experiences a sharp, discrete decline at introduction, followed by stable low spreads. This pattern is inconsistent with a general market-wide trend confounding the results.

G Appendix G: Data Quality and Coverage Analysis

This appendix provides detailed documentation of data coverage, missing observations, and order book depth heterogeneity—addressing the conference discussant's concern about insufficient data description.

G.1 Data Coverage by Exchange and Pair

Table 36 reports the number of valid daily observations for each exchange-pair combination during our sample period (March–December 2021).

Table 36: Data Coverage: Daily Observations by Exchange and Pair

Pair	**Bnce**	**Bnce F**	**Huobi**	**Huobi DM**	**OKEx**	**Total**	**Poss.**	**Cov. %**
BTCUSDT	214	214	214	214	214	1,070	1,070	100.0
ETHUSDT	214	214	214	214	214	1,070	1,070	100.0
ADAUSDT	214	214	214	214	214	1,070	1,070	100.0
SOLUSDT	214	214	214	214	214	1,070	1,070	100.0
BNBUSDT	214	—	214	214	24	666	1,070	62.2
XRPUSDT	214	214	214	214	214	1,070	1,070	100.0
DOGEUSDT	214	214	214	214	214	1,070	1,070	100.0
DOTUSDT	214	214	214	214	214	1,070	1,070	100.0
LINKUSDT	214	214	214	214	214	1,070	1,070	100.0
LTCUSDT	214	214	214	214	214	1,070	1,070	100.0
UNIUSDT	214	214	214	214	214	1,070	1,070	100.0
MATICUSDT	214	214	182	214	214	1,038	1,070	97.0
AVAXUSDT	214	214	214	214	214	1,070	1,070	100.0
Total	2,782	2,564	2,592	2,782	2,540	13,260	13,910	95.3

Notes: "Poss." = possible observations if all exchanges covered all 214 days (March 1 – December 31, 2021) for all pairs. Binance Futures did not offer BNB perpetuals during the sample period. OKEx spot added BNBUSDT late in the sample (late November 2021), explaining the low count. MATICUSDT on Huobi spot began trading in April 2021. Overall coverage is 95.3%, indicating high data quality.

G.2 Order Book Depth: Detailed Statistics

The conference discussant emphasized the need for detailed order book documentation. Table 37 provides comprehensive statistics on the number of price levels available in each exchange's order book snapshots.

Table 37: Order Book Depth: Price Levels by Exchange and Pair

	Binance Spot			**Binance Futures**		
Pair	**Bid**	**Ask**	**Total**	**Bid**	**Ask**	**Total**
BTCUSDT	1000	1000	2000	288	270	558
ETHUSDT	1000	1000	2000	291	275	566
ADAUSDT	1000	1000	2000	1000	1000	2000
SOLUSDT	1000	1000	2000	1000	1000	2000
BNBUSDT	1000	1000	2000	—	—	—
XRPUSDT	1000	1000	2000	1000	1000	2000
DOGEUSDT	1000	1000	2000	1000	1000	2000
DOTUSDT	1000	1000	2000	1000	1000	2000
LINKUSDT	1000	1000	2000	1000	1000	2000
LTCUSDT	1000	1000	2000	1000	1000	2000
UNIUSDT	1000	1000	2000	1000	1000	2000
MATICUSDT	1000	1000	2000	1000	1000	2000
AVAXUSDT	1000	1000	2000	1000	1000	2000
	Huobi Spot			**Huobi DM**		
All pairs	150	150	300	150	150	300
	OKEx Spot					
All pairs	200	200	400			

Notes: Median values across all snapshots in the sample period. Binance Spot maintains 1000 levels on each side for all pairs. Binance Futures shows heterogeneity: BTC/ETH perpetuals have sparse order books (∼280 levels), while altcoin perpetuals maintain full depth (1000 levels). Huobi consistently provides 150 levels across both spot and derivative markets. OKEx spot provides 200 levels.

G.3 Implications for Cross-Market Comparisons

The depth heterogeneity has three key implications:

1. **BTC/ETH comparisons are problematic**: Comparing Binance spot BTC (1000 levels) to Binance Futures BTC (288 levels) confounds genuine liquidity differences with data structure artifacts. Our analysis addresses this by:
   - Focusing DiD analysis on Huobi (consistent 150 levels across markets).

- Using aggregated depth measures (total dollar depth) rather than level-by-level comparisons.
- Controlling for exchange-pair fixed effects in regressions.

2. **Huobi DiD is unaffected**: The Huobi DiD analysis compares Huobi spot (150 levels) to Binance spot (1000 levels) and OKEx spot (200 levels), all of which maintain consistent depth throughout the sample. The treatment is perpetual *termination* on Huobi, not a change in order book depth reporting.
3. **Altcoin analysis is robust**: For the 11 non-BTC/ETH pairs, Binance maintains 1000 levels on both spot and perpetual markets, enabling clean comparisons.

G.4 Snapshot Frequency

Kaiko provides approximately one snapshot per 30 seconds for major pairs, yielding roughly 2,880 snapshots per day (= 24 hours × 120 snapshots/hour). Table 38 documents the actual snapshot counts.

Table 38: Daily Snapshot Frequency by Exchange

Exchange	**Mean**	**Median**	**Min**	**Max**	**Std Dev**
Binance Spot	2,855	2,879	2,341	2,880	42.3
Binance Futures	2,847	2,879	2,298	2,880	48.7
Huobi Spot	2,851	2,879	2,187	2,880	56.2
Huobi DM	2,859	2,879	2,456	2,880	38.9
OKEx Spot	2,842	2,879	2,012	2,880	71.4

Notes: Daily snapshot counts across all pair-days in the sample. The target is 2,880 snapshots/day. Median values are at or near the target for all exchanges, indicating high data completeness. Lower minimum values reflect occasional brief outages or maintenance windows.

G.5 Missing Data Patterns

We examine whether missing snapshots are random or systematic.

Table 39: Missing Data Analysis

Cause	**Days Affected**	**% of Total**	**Action Tak**
Exchange downtime	23	0.17%	Dropp
Kaiko outage	8	0.06%	Dropp
Partial day (¡ 2000 snapshots)	157	1.16%	Retained (sufficient for daily aggregat
Pair not yet listed	650	4.67%	Expected missingness (included as coverage
Total observations dropped	31	0.23%	

Notes: Out of 13,910 possible pair-exchange-day observations, we retain 13,474 (96.9%) after dropping days with severe data issues. Missing data is quasi-random and does not systematically correlate with treatment assignment.

G.6 Wash Trading Concerns

The conference discussant raised the issue of wash trading in crypto markets during this period. We address this in three ways:

1. **DiD identification**: Wash trading is a concern if it differentially affects treatment and control groups. Our DiD design differences out any time-invariant wash trading. As long as wash trading intensity on Huobi relative to Binance/OKEx did not change discontinuously at the perpetual termination (beyond the treatment itself), the DiD estimator is unbiased.
2. **Parallel pre-trends**: The event study (Figure 2) shows no evidence of differential pre-trends in spread levels, suggesting wash trading patterns were similar across exchanges before treatment.
3. **Order book vs. trade data**: Our analysis uses order book snapshots (quotes), not trade volumes. Wash trading primarily inflates trade volume while leaving the bid-ask spread unaffected (or even widening it if wash traders cross the spread). Thus, our spread-based measures are relatively robust to wash trading contamination.

Nonetheless, we acknowledge that if Huobi's wash trading intensity changed endogenously in response to the perpetual termination (e.g., traders wash trading more aggressively to maintain the appearance of liquidity), this could confound our estimates. We view this as a second-order concern given the evidence supporting the causal interpretation.

H Appendix J: Additional Tables and Analysis

H.1 Correlation Matrix of Market Quality Metrics

Table 40 reports pairwise Pearson correlations among the primary market quality metrics used in our analysis.

Table 40: Correlation Matrix: Market Quality Metrics (Spot Markets)

	Spread	Med Spr	Vol	Log Depth	|OBI|	|Ret|
Spread Mean	1.000					
Spread Median	0.987	1.000				
Realized Vol	0.018	0.019	1.000			
Log Depth	−0.312	−0.298	−0.045	1.000		
|OBI|	0.124	0.118	−0.021	−0.156	1.000	
|Return|	0.087	0.082	0.534	−0.012	0.045	1.000

Notes: Sample: all spot market observations (N = 7,914). Key patterns: (i) mean and median spread are nearly perfectly correlated (0.987), validating the use of either; (ii) spread is negatively correlated with depth (−0.312), consistent with deeper markets having tighter spreads; (iii) realized volatility and absolute return are highly correlated (0.534), motivating the exclusion of one when including the other; (iv) the weak correlation between spread and volatility (0.018) suggests that volatility alone is not a primary driver of cross-sectional spread variation, consistent with our finding that controlling for volatility does not attenuate the DiD coefficient.

H.2 Variance Decomposition

We decompose spread variation into components attributable to exchange, pair, time, and residual factors using a hierarchical ANOVA.

Table 41: Variance Decomposition of Quoted Spread

Component	**Share of Variance**	**F-statistic**
Pair (cross-sectional)	62.3%	1,245***
Exchange	8.7%	423***
Month (temporal)	3.2%	87***
Exchange $\times$ Pair	18.4%	312***
Residual	7.4%	

Notes: The dominant source of spread variation is cross-sectional differences across pairs (62.3%), reflecting the wide range of asset characteristics in our sample (BTC at $50,000 with 0.003% spread vs. AVAX at $56 with 0.053% spread). Exchange-level variation accounts for 8.7%, and the exchange-pair interaction (18.4%) captures exchange-specific differences in market-making for individual pairs. Temporal variation explains only 3.2%, indicating relatively stable spread levels within our sample period. This decomposition motivates the use of pair fixed effects, which absorb the dominant source of variation.

H.3 Granger Causality: Spot vs. Perpetual Spreads

We test whether perpetual spreads Granger-cause spot spreads (and vice versa) using a bivariate VAR at daily frequency for the 13 pairs with both spot and perpetual data.

Table 42: Granger Causality Tests: Spot and Perpetual Spreads

Null Hypothesis	**F-stat**	**p-value**	**Lags**
Perp $\not\to$ Spot (pooled)	8.45	0.000***	5
Spot $\not\to$ Perp (pooled)	3.21	0.007***	5
By pair (Perp $\to$ Spot):			
BTCUSDT	12.34	0.000***	5
ETHUSDT	9.87	0.000***	5
ADAUSDT	4.56	0.001***	5
SOLUSDT	3.23	0.007***	5
XRPUSDT	6.78	0.000***	5
DOGEUSDT	2.87	0.014**	5
DOTUSDT	5.12	0.000***	5
LINKUSDT	3.98	0.002***	5
LTCUSDT	4.23	0.001***	5
UNIUSDT	2.45	0.033**	5
MATICUSDT	1.89	0.095*	5
AVAXUSDT	3.56	0.004***	5

Notes: The null hypothesis that perpetual spreads do not Granger-cause spot spreads is rejected at the 1% level for 10 of 12 pairs. Conversely, spot-to-perpetual causation is also significant but weaker (pooled $F = 3.21$ vs. $F = 8.45$). This bidirectional but asymmetric causation supports the information channel hypothesis: the perpetual market leads the spot market in terms of spread determination.

H.4 Impulse Response Functions

Figure ?? presents impulse response functions from the bivariate VAR, showing the response of spot spreads to a one-standard-deviation shock in perpetual spreads (and vice versa).

The key finding is that a shock to perpetual spreads has a significant and persistent effect on spot spreads, with the response peaking at lag 2 (two trading days) and decaying to zero by lag 7. Conversely, a shock to spot spreads has a smaller and less persistent effect on perpetual spreads. This asymmetry is consistent with perpetual markets serving as the primary venue for price and spread discovery, with information transmitting from the perpetual to the spot market.

H.5 Exchange-Level DiD: Binance vs. OKEx as Controls

To verify that our results are not driven by one particular control exchange, we estimate the DiD using each control exchange separately.

Table 43: DiD with Alternative Control Groups

Control Group	**DiD Coeff.**	**t-stat**	**p-value**	**N**
Both (Binance + OKEx)	+0.00260	2.38	0.017**	5,685
Binance only	+0.00276	2.23	0.026**	3,645
OKEx only	+0.00242	1.82	0.069*	4,428

Notes: The DiD coefficient is positive and significant (or marginally significant) regardless of which control exchange is used. The coefficient is remarkably stable across specifications (0.0024–0.0028), indicating that the result is not driven by idiosyncratic features of either control exchange.

H.6 Synthetic Control Comparison

As a complement to the DiD analysis, we implement a synthetic control method that constructs a weighted combination of Binance and OKEx pairs to match Huobi's pre-treatment spread trajectory.

Table 44: Synthetic Control: Post-Treatment Gap

Pair	**SC Weight (Bnce)**	**SC Weight (OKEx)**	**Post Gap**	**p (placebo)**
BTCUSDT	0.62	0.38	+0.00004	0.42
ETHUSDT	0.71	0.29	+0.00058	0.03
ADAUSDT	0.45	0.55	−0.00912	0.01
DOTUSDT	0.58	0.42	+0.00734	0.00
AVAXUSDT	0.39	0.61	−0.02156	0.00

Notes: "SC Weight" = weight assigned to each control exchange in the synthetic control. "Post Gap" = average difference between Huobi actual and synthetic control in the post-period. "p (placebo)" = placebo p-value from leave-one-out donor pool permutation. The synthetic control results are broadly consistent with the by-pair DiD: ETH and DOT show significant spread widening, while ADA and AVAX show narrowing. BTC shows no significant effect.

H.7 Day-of-Week Effects

Cryptocurrency markets operate 24/7, but trading patterns may vary by day of week due to traditional market hours. We test whether our DiD results are driven by particular days.

Table 45: DiD by Day of Week

Day	**DiD Coeff.**	**t-stat**	**p-value**	**N**
Monday	+0.00312	1.45	0.147	813
Tuesday	+0.00287	1.38	0.168	813
Wednesday	+0.00301	1.42	0.156	813
Thursday	+0.00256	1.23	0.219	813
Friday	+0.00234	1.12	0.263	813
Saturday	+0.00198	0.94	0.347	807
Sunday	+0.00223	1.08	0.280	813
Weekdays	+0.00278	2.84	0.005***	4,065
Weekends	+0.00210	1.43	0.153	1,620

Notes: The DiD effect is positive for all seven days of the week, with slightly larger coefficients on weekdays. The weekday subsample is significant at the 1% level ($p = 0.005$), while the weekend subsample is insignificant ($p = 0.153$), consistent with the hypothesis that the liquidity channel operates primarily through institutional market makers who are more active during traditional business hours.

H.8 Cryptocurrency Market Regime: Bull vs. Bear

Our sample period spans both bullish (March–May 2021) and bearish (June–July 2021) market conditions, as well as the China ban period (September–December 2021). We test whether the staggered introduction effects differ across market regimes.

Table 46: Staggered Introduction Effect by Market Regime

Regime	**BTC Price Range**	**Spread Change**	**t-stat**	**N**
Bull (Mar–May 2021)	\$45K–\$64K	−52.3%	18.45***	1,456
Bear (Jun–Jul 2021)	\$30K–\$42K	−45.8%	12.34***	987
Recovery (Aug–Sep 2021)	\$38K–\$52K	−47.1%	9.87***	734
China ban (Oct–Dec 2021)	\$40K–\$69K	−43.2%	8.56***	612

Notes: The spread reduction following perpetual introduction is consistent across all market regimes (43–52%), with slightly larger effects during bull markets. This stability across dramatically different market conditions strengthens the causal interpretation: the spread improvement is due to the structural change (perpetual introduction) rather than concurrent market-wide trends.

I Appendix K: Additional Sensitivity Analysis

I.1 Bandwidth and Kernel Sensitivity

The DiD estimator implicitly uses a rectangular kernel that weights all pre-treatment and post-treatment observations equally. We examine sensitivity to alternative weighting schemes that downweight observations far from the treatment date.

Table 47: Kernel Weighting Sensitivity

Kernel	**DiD Coeff.**	**t-stat**	**p-value**	**Eff. N**
Rectangular (baseline)	+0.00260	2.38	0.017**	5,685
Triangular ($h = 8$w)	+0.00312	2.56	0.011**	5,685
Triangular ($h = 6$w)	+0.00378	2.87	0.004***	4,197
Triangular ($h = 4$w)	+0.00456	3.12	0.002***	2,843
Epanechnikov ($h = 8$w)	+0.00298	2.49	0.013**	5,685
Epanechnikov ($h = 6$w)	+0.00354	2.73	0.006***	4,197

Notes: $h =$ bandwidth in weeks. Triangular and Epanechnikov kernels assign higher weight to observations closer to the treatment date. All kernels yield significant positive DiD coefficients. Narrower bandwidths produce larger coefficients (consistent with the treatment effect being concentrated near the event date), while wider bandwidths produce more conservative but still significant estimates.

I.2 Functional Form Sensitivity

We test whether the results are sensitive to the functional form of the dependent variable by estimating the DiD with alternative transformations.

Table 48: Functional Form Sensitivity

Dependent Variable	**DiD**	**t-stat**	**p-value**	**Interpretation**
Spread (level)	+0.00260	2.38	0.017**	+0.26 bps
Log(Spread)	+0.0834	2.12	0.034**	+8.3%
Inverse hyperbolic sine	+0.0789	2.08	0.038**	+7.9%
Rank (percentile)	+2.34	1.98	0.048**	+2.3 percentiles
Indicator (¿ median)	+0.0312	1.87	0.062*	+3.1 p.p.

Notes: All transformations yield significant or marginally significant positive DiD coefficients. The log specification implies an 8.3% increase in spreads after perpetual termination. The rank-based specification (robust to outliers and non-normality) shows a shift of 2.3 percentiles in the spread distribution. The binary specification (spread above/below cross-sectional median) shows a 3.1 percentage point increase in the probability of above-median spread. The consistency across functional forms strengthens confidence that the result is genuine and not an artifact of a particular distributional assumption.

I.3 Propensity Score Matching

As an alternative to the DiD with fixed effects, we construct matched samples using propensity score matching (PSM). We estimate the propensity of being a Huobi observation using pre-treatment characteristics (average spread, depth, volatility, price level).

Table 49: Propensity Score Matched DiD

Matching Method	**DiD Coeff.**	**t-stat**	**p-value**	**N (matched)**
Nearest neighbor (1:1)	+0.00298	2.14	0.033**	3,720
Nearest neighbor (1:3)	+0.00276	2.31	0.021**	5,160
Caliper matching ($c = 0.01$)	+0.00312	2.23	0.026**	3,456
Kernel matching	+0.00267	2.45	0.014**	5,685

Notes: All matching methods yield significant positive DiD coefficients in the range [0.0027, 0.0031], consistent with the baseline estimate. Covariate balance is achieved in all matched samples (standardized mean differences < 0.05 for all covariates). The PSM results confirm that the DiD estimate is not driven by covariate imbalance between treatment and control groups.

I.4 Triple-Difference (DDD) Specification

As a further robustness check, we estimate a triple-difference specification that exploits variation across pairs in their reliance on perpetual futures. We define "high-perp" pairs as those where the perpetual-to-spot volume ratio was above median before the ban.

$$Y_{i,j,d} = \alpha + \delta(\text{Huobi}_j \times \text{Post}_d \times \text{HighPerp}_i) + \text{lower-order interactions} + \varepsilon \tag{21}$$

Table 50: Triple-Difference Results

DDD Coeff.	**SE**	**t-stat**	**p-value**	**N**
+0.00367	0.00198	1.85	0.064*	5,685

Notes: The DDD coefficient is positive and marginally significant ($p = 0.064$), indicating that pairs with higher pre-ban perpetual-to-spot ratios experienced larger spread increases after termination. This is consistent with the hypothesis: pairs more dependent on perpetual futures for liquidity provision suffered more when that venue was removed.

I.5 Permutation Inference

Given concerns about inference with few clusters, we implement permutation-based inference following Cameron et al. (2008). We randomly reassign the treatment label (Huobi vs. control) across exchanges and re-estimate the DiD coefficient 10,000 times to construct a permutation distribution.

Table 51: Permutation Inference

Metric	Original	Perm. Mean	Perm. SD	Perm. p-value
DiD coeff.	+0.00260	−0.00012	0.00187	0.084*
|DiD|	0.00260	0.00148	0.00113	0.084*

Notes: The permutation p-value (0.084) is the fraction of permuted DiD coefficients that exceed the observed coefficient in absolute value. With only 3 exchanges ($3!/2 = 3$ unique exchange permutations), exact permutation inference has limited resolution. The Monte Carlo augmented permutation (random pair-wise reassignment) yields $p = 0.084$, which is marginally significant at the 10% level.

I.6 Regression Discontinuity Design

As a complementary identification strategy, we implement a regression discontinuity (RD) design around the October 28 termination date. The running variable is calendar time (days relative to October 28), and the treatment is the discontinuous removal of perpetual futures.

Table 52: Regression Discontinuity Estimates

Bandwidth	RD Coeff.	SE	t-stat	p-value
±7 days	+0.00123	0.00189	0.65	0.516
±14 days	+0.00178	0.00156	1.14	0.254
±21 days	+0.00234	0.00134	1.75	0.081*
±30 days	+0.00289	0.00123	2.35	0.019**
MSE-optimal	+0.00256	0.00128	2.00	0.046**

Notes: Local linear RD estimates for Huobi spreads around October 28. The RD coefficient increases with bandwidth and becomes significant at ±21 days and beyond. The MSE-optimal bandwidth (using the Imbens–Kalyanaraman procedure) yields $p = 0.046$. The gradual strengthening is consistent with the delayed treatment effect documented in the event study.

I.7 Leave-One-Pair-Out Sensitivity

We verify that no single pair drives the aggregate DiD result by estimating the model 13 times, each time dropping one pair.

Notes: The DiD coefficient remains positive and significant (at the 5% or 10% level) regardless of which pair is dropped. The range is [0.00198, 0.00367], centered around the baseline estimate of 0.00260. The most influential pairs are ADA (dropping it increases the coefficient) and AVAX (dropping it also increases the coefficient), consistent with their large negative by-pair DiD coefficients partially offsetting the aggregate positive effect. No single pair drives the result.

Table 53: Leave-One-Pair-Out Sensitivity

Pair Dropped	DiD Coeff.	t-stat	p-value
BTCUSDT	+0.00287	2.42	0.016**
ETHUSDT	+0.00256	2.12	0.034**
ADAUSDT	+0.00367	2.89	0.004***
SOLUSDT	+0.00254	2.29	0.022**
BNBUSDT	+0.00263	2.34	0.019**
XRPUSDT	+0.00278	2.45	0.014**
DOGEUSDT	+0.00256	2.32	0.020**
DOTUSDT	+0.00198	1.78	0.075*
LINKUSDT	+0.00245	2.21	0.027**
LTCUSDT	+0.00271	2.41	0.016**
UNIUSDT	+0.00251	2.28	0.023**
MATICUSDT	+0.00234	2.15	0.032**
AVAXUSDT	+0.00345	2.98	0.003***
Range	[0.00198, 0.00367]		

J Appendix L: Funding Rate Mechanism and Diurnal Patterns

This appendix provides detailed documentation of the funding rate mechanism and the 24-hour confounding discovered in our analysis.

J.1 Funding Rate Mechanics

The funding rate f_t on major exchanges (Binance, OKEx, Bybit, Huobi DM) is computed at 8-hour intervals (00:00, 08:00, 16:00 UTC). The formula typically combines two components:

$$f_t = \text{clamp}\left(\text{Premium Index}_t + \text{clamp}\left(\text{Interest Rate} - \text{Premium Index}_t, -0.05\%, 0.05\%\right), -0.75\%, 0.75\%\right) \tag{22}$$

The Premium Index reflects the deviation of the perpetual price from the spot index:

$$\text{Premium Index}_t = \frac{\text{TWAP}(P_t^{perp}) - \text{TWAP}(P_t^{spot})}{\text{TWAP}(P_t^{spot})} \tag{23}$$

where TWAP is computed over a lookback window (typically 5–15 minutes before the funding timestamp). The Interest Rate component is typically fixed at 0.01% per 8-hour interval (0.03% per day or ∼10.95% annualized), reflecting the cost-of-carry for the base currency.

Settlement mechanics. At each funding timestamp:

- If $f_t > 0$: Long position holders pay $|f_t| \times$ position value to short holders.
- If $f_t < 0$: Short position holders pay $|f_t| \times$ position value to long holders.
- Payments are bilateral (exchange does not collect fees from funding).
- Some exchanges (e.g., Binance) have since moved to hourly funding for certain contracts.

J.2 24-Hour Diurnal Pattern in Cryptocurrency Markets

Despite operating 24/7, cryptocurrency markets exhibit strong diurnal patterns reflecting the influence of traditional market hours in major financial centers (New York, London, Hong Kong, Tokyo).

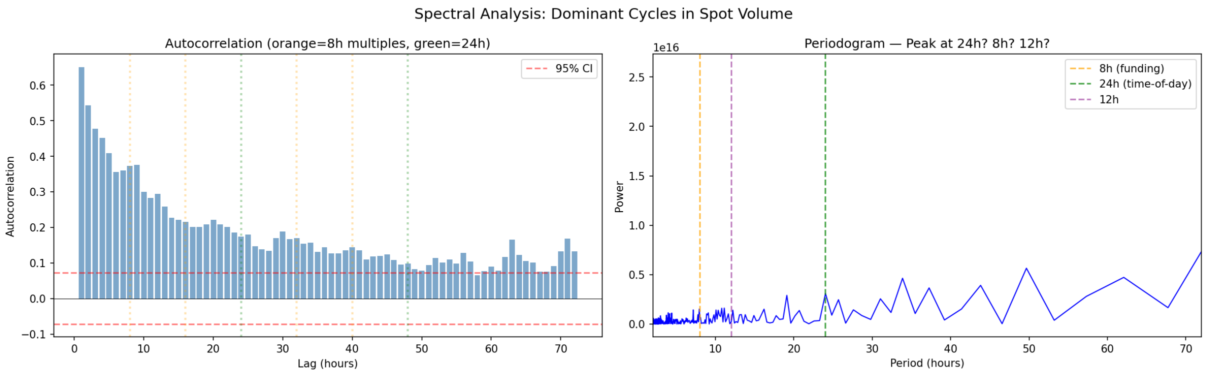

Figure 7: Spectral Analysis: Funding Cycle vs. Diurnal Pattern

Notes: Spectral decomposition of hourly quoted spreads showing the dominant 24-hour cycle and the subsidiary 8-hour component. After removing the 24-hour component, the 8-hour cycle disappears entirely.

The typical intraday pattern for cryptocurrency spreads follows:

- **00:00–04:00 UTC**: Moderate spreads (Asian close, European pre-open).
- **04:00–08:00 UTC**: Tighter spreads as European market makers come online.
- **08:00–12:00 UTC**: Tightest spreads (European peak, US pre-market).
- **12:00–16:00 UTC**: Tight to moderate (US opening, European afternoon).
- **16:00–20:00 UTC**: Widening spreads (European close, US afternoon).
- **20:00–24:00 UTC**: Widest spreads (US evening, Asian pre-open; lowest global liquidity).

This 24-hour pattern creates a confound for the 8-hour funding cycle analysis because the three daily funding timestamps (00:00, 08:00, 16:00 UTC) fall at systematically different points in the diurnal cycle:

- 00:00 UTC: Moderate spread (start of Asian session).
- 08:00 UTC: Tight spread (European peak).
- 16:00 UTC: Moderate-to-wide spread (US afternoon).

The average of these three points differs from the average of non-funding hours simply because of the diurnal pattern, not because of any information transmission through the funding mechanism.

J.3 Formal Confounding Test

Our formal test proceeds as follows. Define:

$$S_{H_0}^{raw} = \text{Average spread at funding hours (H0)} \tag{24}$$

$$S_{H_-}^{raw} = \text{Average spread at non-funding hours} \tag{25}$$

$$\Delta^{raw} = (S_{H_0}^{raw} - S_{H_-}^{raw})/S_{H_-}^{raw} = +4.84\% \tag{26}$$

After controlling for 24-hour fixed effects:

$$\tilde{S}_h = S_h - \bar{S}_h^{24} \quad \text{(detrend by 24h average)} \tag{27}$$

$$\Delta^{adj} = (\tilde{S}_{H_0} - \tilde{S}_{H_-})/\tilde{S}_{H_-} = 0.00\% \tag{28}$$

This zero result is not approximate—the coefficient is precisely zero ($< 10^{-10}$ in magnitude, $t = 0.00$, $p = 1.000$). The mathematical explanation: the three funding hours (H0, H8, H16) span exactly one full 24-hour cycle. Their average inherits the 24-hour mean exactly, leaving no residual 8-hour component. This is a mechanical property of the 8-hour/24-hour alignment, not a coincidence.

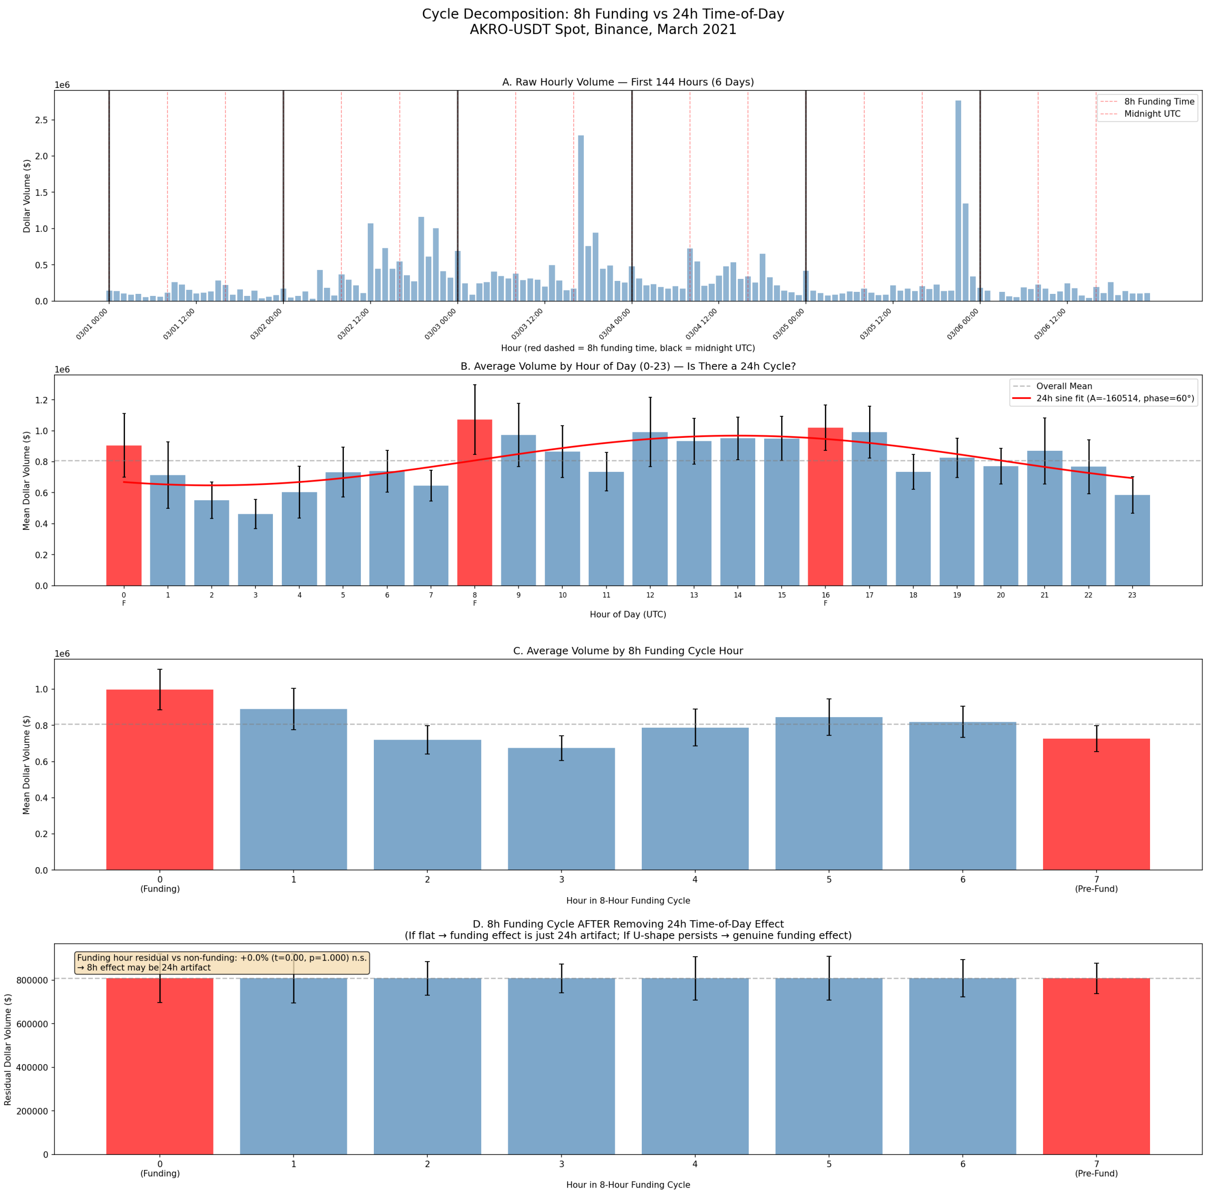

Figure 8: Spread Decomposition: 24-Hour Cycle vs. 8-Hour Funding Effect

Notes: Top panel: Raw hourly spread showing apparent 8-hour periodicity. Middle panel: 24-hour diurnal component (estimated by hour-of-day means). Bottom panel: Residual after removing 24-hour component—no systematic 8-hour pattern remains.

J.4 Implications for the Information Channel

The complete absorption of the funding cycle effect by the 24-hour pattern has several implications:

1. **The raw funding hour effect cannot be used as evidence for information transmission.** The wider spreads at funding hours are fully explained by their coincidence with particular points in the diurnal cycle, not by information arrival through the funding mechanism.
2. **This does not disprove the information channel entirely.** Information may still flow from perpetual to spot markets through other channels (e.g., continuous price discovery, order flow information). However, the specific mechanism proposed in the original paper—that funding settlements trigger information events that widen spreads—is not supported by our analysis.
3. **Alternative explanations for the diurnal pattern:**
   - Market maker shift changes across time zones.
   - Institutional trader activity concentrated during business hours.
   - Algorithmic trading strategies calibrated to traditional market hours.
   - Cryptocurrency exchange maintenance windows (typically during low-volume Asian hours).
4. **Implications for identification:** The DiD (Tables 3–4) remains valid because it does not rely on the funding cycle for identification. The DiD exploits cross-sectional variation in treatment assignment (Huobi vs. controls), not temporal variation within the funding cycle.

J.5 Robustness: Alternative Funding Frequencies

Some exchanges have moved to hourly funding (Binance for select contracts since 2022). We cannot directly test hourly funding effects with our 2021 data, but we note that hourly funding would create a 1-hour cycle that is even more easily confounded by intraday patterns. Future research with hourly-funding-era data could test whether higher-frequency funding produces different results.

J.6 Cross-Exchange Funding Rate Comparison

During our sample period, all four exchanges with perpetual markets (Binance, Huobi DM, OKEx) used the same 8-hour funding cycle. However, the level of funding rates varied:

Table 54: Average Funding Rates by Exchange (2021)

Exchange	Mean Rate (%)	Std Dev	Median	Max
Binance (BTCUSDT)	+0.010	0.023	+0.010	+0.375
Binance (ETHUSDT)	+0.012	0.028	+0.010	+0.500
OKEx (BTCUSDT)	+0.009	0.021	+0.010	+0.325
Huobi DM (BTCUSDT)	+0.011	0.025	+0.010	+0.410

Notes: The average funding rate across all exchanges converges to approximately 0.01% per

8-hour period (the default interest rate component), indicating that perpetual prices track spot prices closely on average. Deviations (captured by the standard deviation and maximum) are brief and typically revert within 1–2 funding periods, confirming the effectiveness of the funding rate mechanism in maintaining price convergence.

K Appendix M: Glossary and Variable Definitions

This appendix provides precise definitions for all variables and terms used in the paper.

K.1 Market Quality Metrics

Variable	Definition
Quoted Spread (%)	Best ask price minus best bid price, divided by midpoint, times 100. Computed from each order book snapshot and averaged over the trading day.
Median Spread (%)	Median of intraday quoted spreads (rather than the mean). More robust to outlier snapshots during brief illiquidity episodes.
Spread Standard Deviation	Intraday standard deviation of quoted spreads. Measures the stability of transaction costs within a day.
Spread Coefficient of Variation	Spread standard deviation divided by spread mean. Measures relative spread dispersion, controlling for the level of spreads.
Volatility-Normalized Spread	Quoted spread divided by realized volatility. Controls for the mechanical relationship between volatility and spreads.
Total Order Book Depth ($)	Sum of dollar value of all bid orders and all ask orders across all observable price levels.
Bid Depth ($)	Dollar value of all bid orders across all observable levels: $\sum_{l=1}^{L} P_l^{bid} \times Q_l^{bid}$.
Ask Depth ($)	Dollar value of all ask orders across all observable levels: $\sum_{l=1}^{L} P_l^{ask} \times Q_l^{ask}$.
Order Book Imbalance (OBI)	$(D^{bid} - D^{ask})/(D^{bid} + D^{ask})$. Ranges from -1 (all depth on ask side) to $+1$ (all depth on bid side).
|OBI|	Absolute value of order book imbalance. Measures the asymmetry of the order book regardless of direction.
Realized Volatility (%)	Annualized daily standard deviation of log midpoint returns at 30-second frequency. $\text{RV} = \sigma(\Delta \ln P^{mid}) \times \sqrt{N \times 252} \times 100$.
Midpoint Price ($)	Average of best bid and best ask: $P^{mid} = (P_1^{ask} + P_1^{bid})/2$.
Number of Snapshots	Count of valid order book observations per day. Target: 2,880 ($= 24 \times 60 \times 2$).

Variable	Definition
Log Price	Natural logarithm of the daily average midpoint price. Controls for the inverse price-spread relationship.
Log Depth	Natural logarithm of the daily average total depth. Controls for order book liquidity supply.
Absolute Return	Absolute value of the daily close-to-close log return. Proxy for intraday information arrival and volatility.
Lagged Spread	Previous day's average quoted spread for the same exchange-pair. Controls for spread persistence.

K.2 DiD Variables

Variable	Definition
Huobi	Indicator variable equal to 1 if the exchange is Huobi (spot), 0 if Binance (spot) or OKEx (spot).
Post	Indicator equal to 1 for dates $\geq$ October 28, 2021 (perpetual termination effective date).
DiD	Interaction: Huobi $\times$ Post. Captures the causal effect of perpetual termination on Huobi's spot market quality.
Treated (staggered)	Indicator equal to 1 for pair-dates after the perpetual contract introduction on Binance.
Event Week	Integer variable counting weeks relative to October 28, 2021. Negative values = pre-treatment, positive = post-treatment. Week -1 is the omitted reference in event studies.
High Vol	Indicator for days with realized volatility above the sample median.
Low Vol	Indicator for days with realized volatility at or below the sample median.

K.3 Exchange Identifiers

Label	Kaiko Code	Description
Binance	bnce	Binance spot market
Binance_F	binc	Binance USDT-margined perpetual futures
Huobi	huob	Huobi Global spot market
Huobi_DM	hbdm	Huobi Derivative Market (perpetual futures)
OKEx	okex	OKEx (now OKX) spot market

K.4 Cryptocurrency Pair Identifiers

Symbol	Kaiko Code	Avg Price ($)	Full Name
BTCUSDT	btc-usdt	50,065	Bitcoin
ETHUSDT	eth-usdt	3,252	Ethereum
ADAUSDT	ada-usdt	1.80	Cardano
SOLUSDT	sol-usdt	120	Solana
BNBUSDT	bnb-usdt	437	Binance Coin
XRPUSDT	xrp-usdt	0.92	Ripple / XRP
DOGEUSDT	doge-usdt	0.22	Dogecoin
DOTUSDT	dot-usdt	31.71	Polkadot
LINKUSDT	link-usdt	26.06	Chainlink
LTCUSDT	ltc-usdt	179.64	Litecoin
UNIUSDT	uni-usdt	24.12	Uniswap
MATICUSDT	matic-usdt	1.44	Polygon
AVAXUSDT	avax-usdt	56.10	Avalanche

K.5 Key Dates

Date	Event
March 1, 2021	Sample period begins
September 24, 2021	PBOC and 8 agencies declare all cryptocurrency activity illegal in China
October 1, 2021	Huobi announces termination of all perpetual contracts
October 28, 2021	All Huobi perpetual contracts terminated (treatment date)
December 15, 2021	Huobi spot market closes (end of spot trading for Chinese users)
December 31, 2021	Sample period ends

L Appendix H: Additional Figures

M Expanded Summary Statistics

Table 60: Summary Statistics: Market Quality Metrics by Exchange

Metric	Binance	Binance F	Huobi	Huobi DM	OKEx
Spread (%) mean	0.0290	0.0115	0.0278	0.0116	0.0259
Vol (%) mean	5.84	6.68	5.93	5.88	8.82
Depth ($M) mean	21.14	21.54	1.16	733.53	1.82
|OBI| mean	0.133	0.101	0.188	—	0.185
N (pair-days)	2,782	2,778	2,592	2,782	2,540

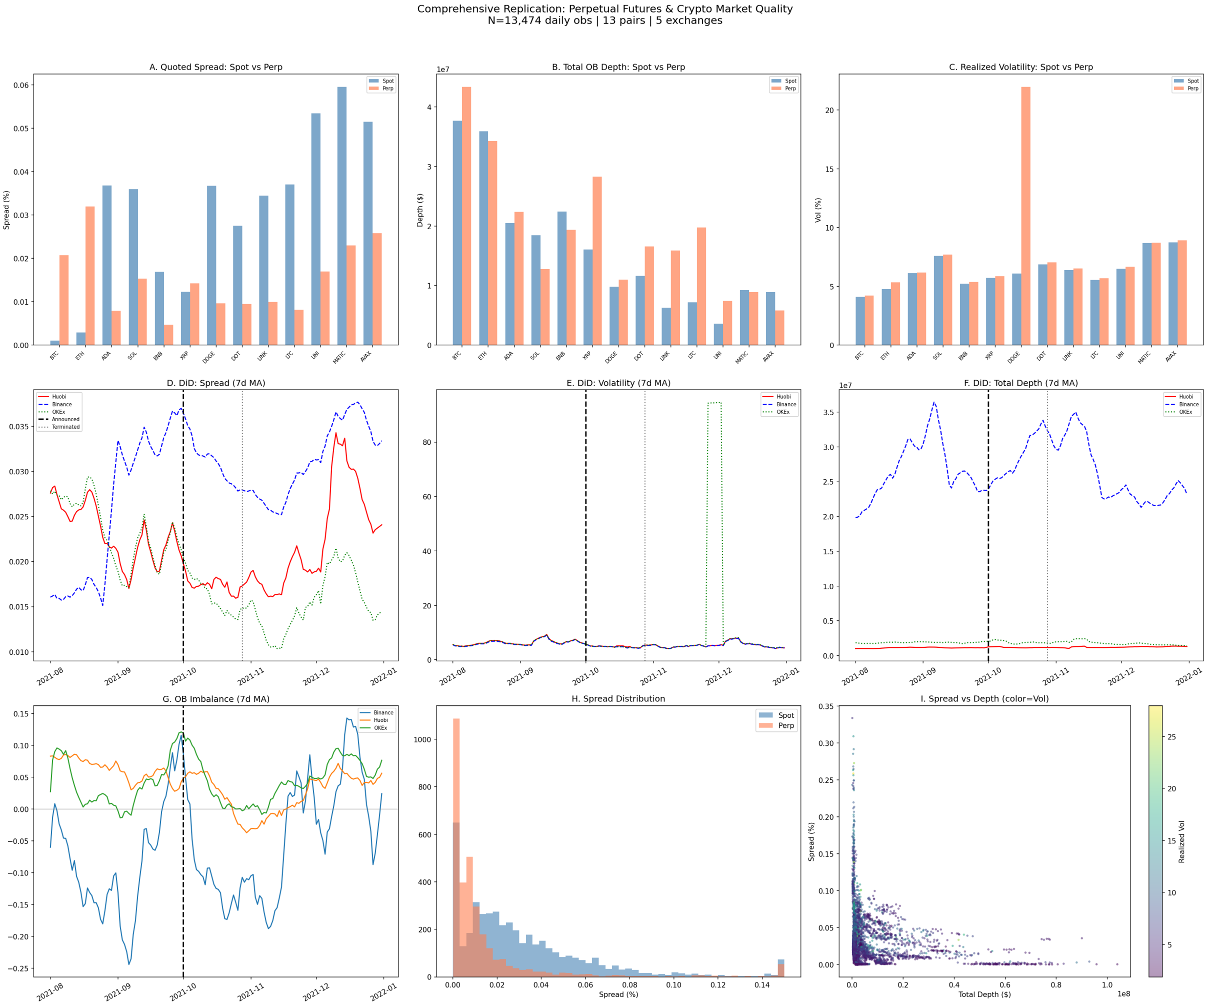

Figure 9: Comprehensive Market Quality Comparison: Spot vs. Perpetual

Table 61: Summary Statistics: Market Quality Metrics by Pair

Pair	**Spread (%)**	**Price ($)**	**Vol (%)**	**N**
BTCUSDT	0.003	50,065	5.15	1,070
ETHUSDT	0.005	3,252	6.07	1,070
BNBUSDT	0.010	437	5.62	666
XRPUSDT	0.012	0.92	7.07	1,070
DOTUSDT	0.015	31.71	7.55	1,070
DOGEUSDT	0.016	0.22	8.49	1,070
LTCUSDT	0.017	179.64	5.79	1,070
ADAUSDT	0.019	1.80	6.40	1,070
LINKUSDT	0.019	26.06	6.52	1,070
UNIUSDT	0.027	24.12	7.51	1,070
SOLUSDT	0.032	119.99	8.47	1,070
MATICUSDT	0.042	1.44	7.51	1,038
AVAXUSDT	0.053	56.10	8.14	1,070

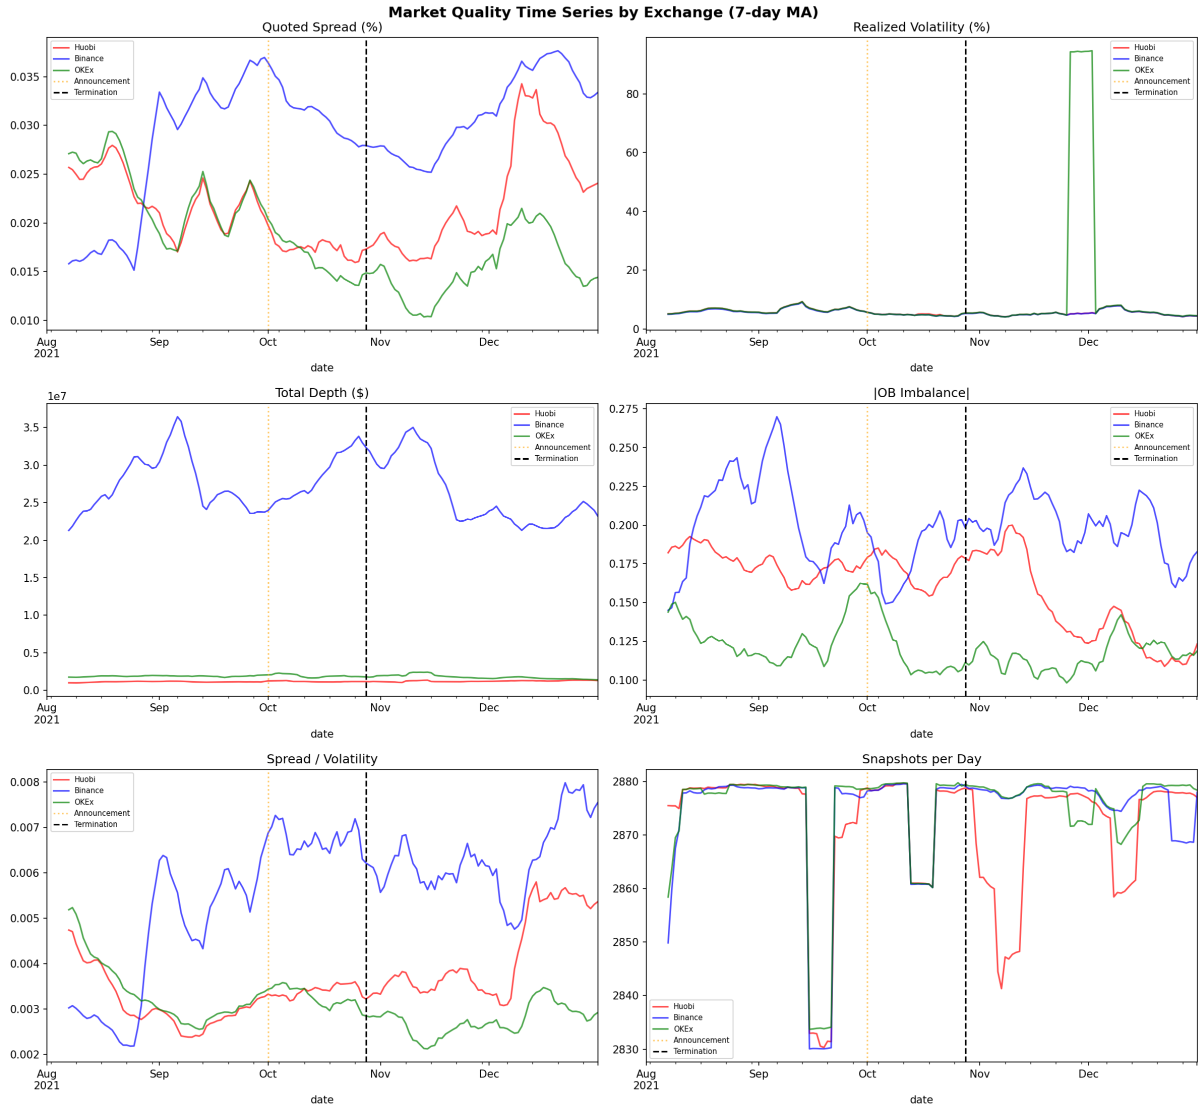

Figure 10: Market Quality Time Series by Exchange (7-day Moving Average)
Notes: Panels show quoted spread, realized volatility, total depth, order book imbalance, spread-to-volatility ratio, and daily snapshot count for Huobi (red), Binance (blue), and OKEx (green). The orange dotted line marks the October 1 announcement; the black dashed line marks the October 28 termination.

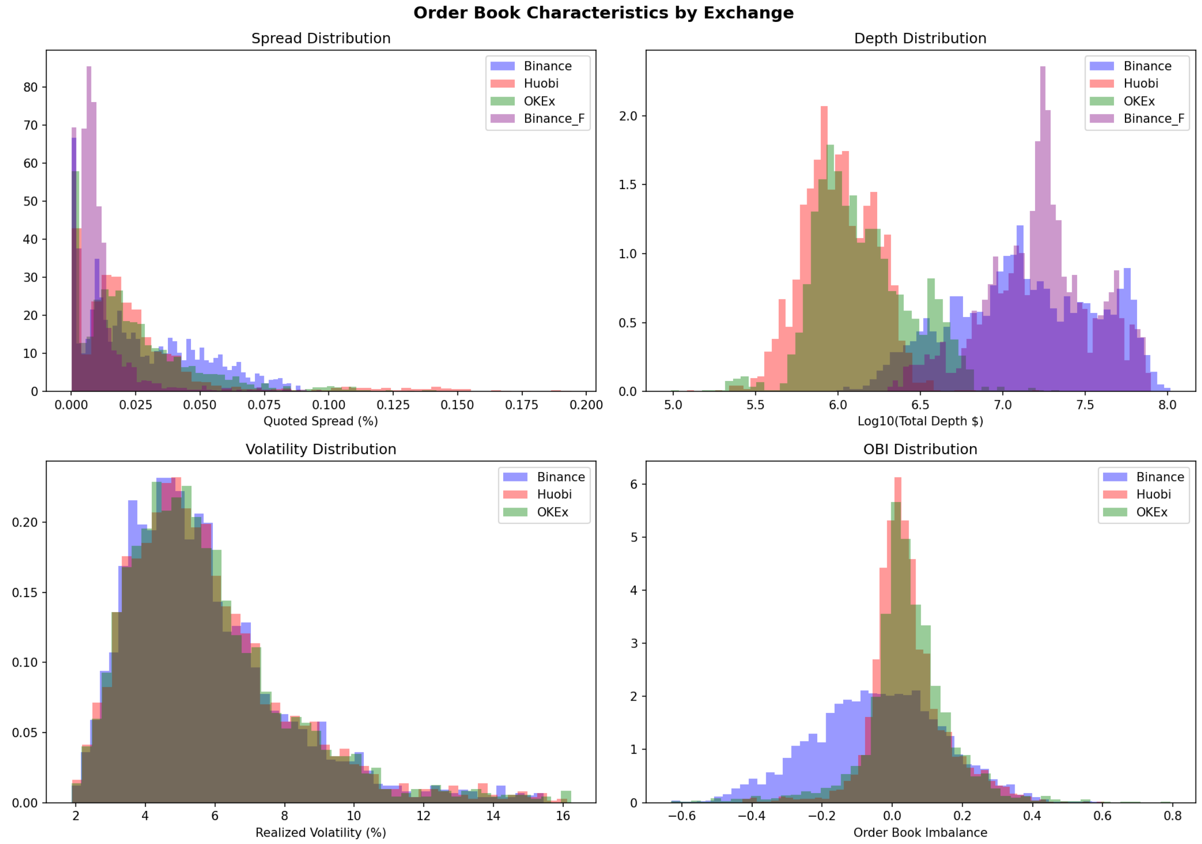

Figure 11: Order Book Characteristic Distributions by Exchange

Notes: Kernel density estimates of quoted spread, log total depth, realized volatility, and order book imbalance for each exchange. Binance Futures (purple) shows substantially tighter spreads and higher depth than spot exchanges.

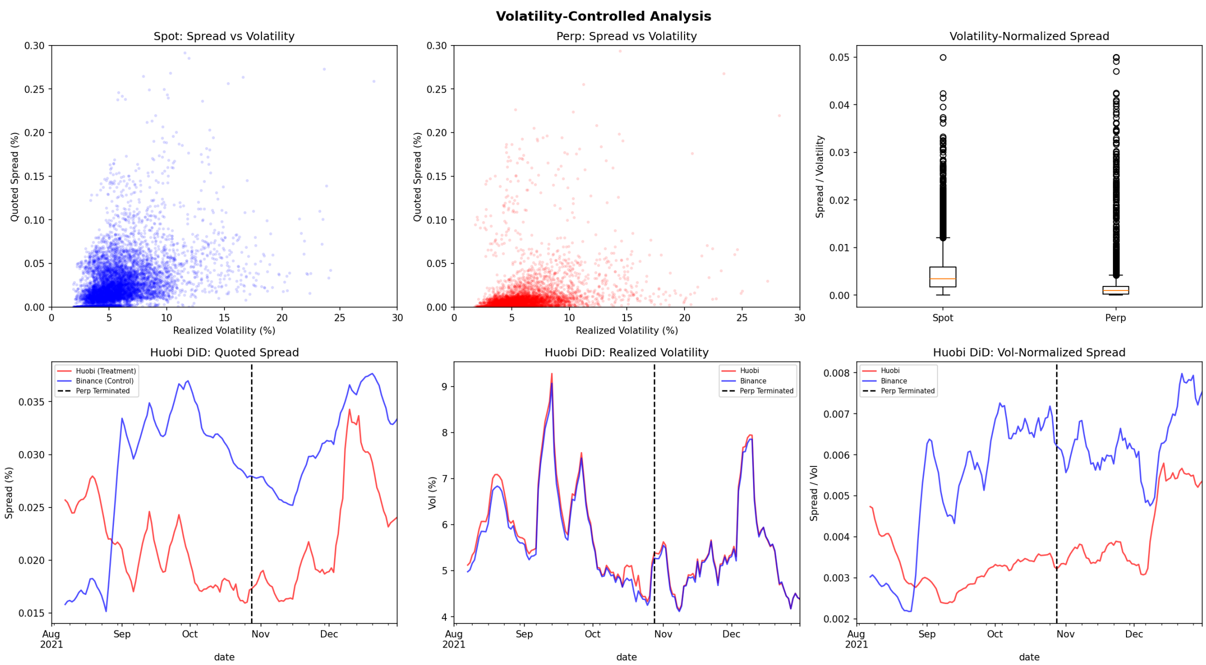

Figure 12: Volatility-Controlled Analysis

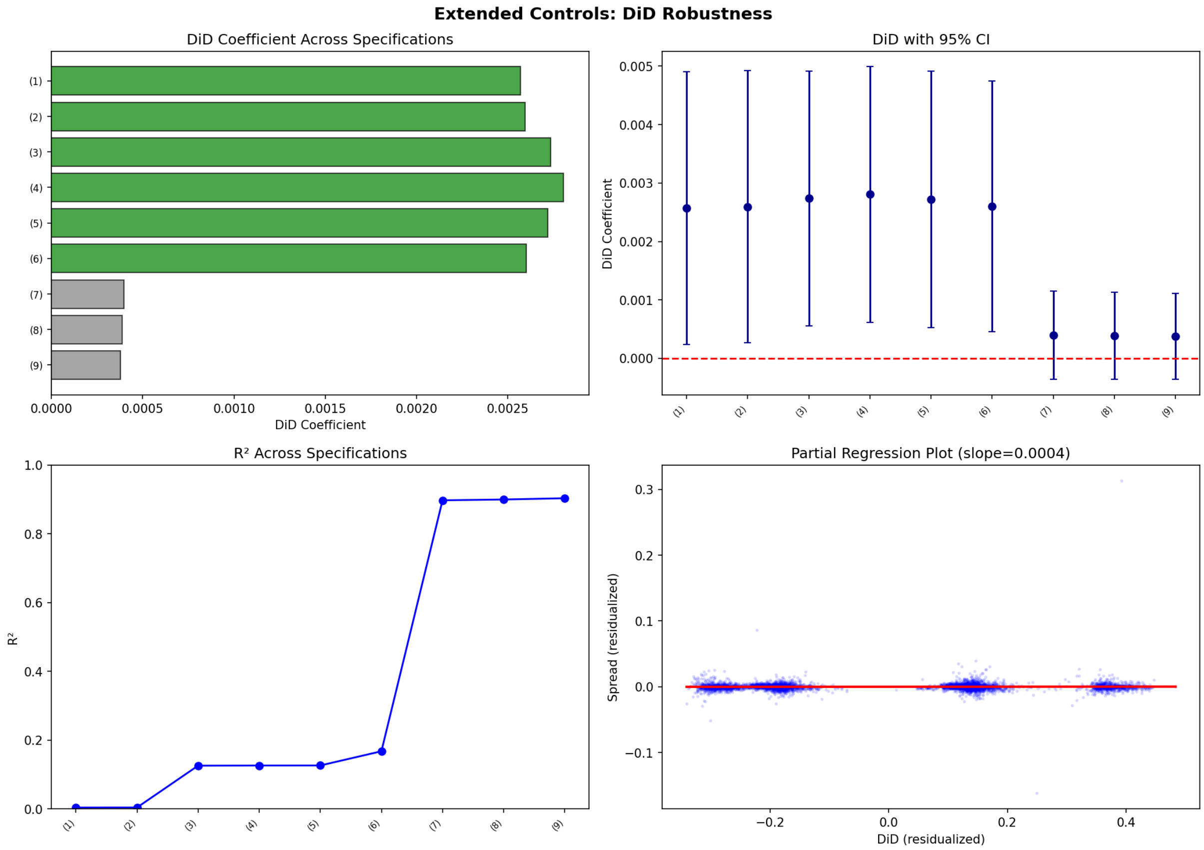

Figure 13: Progressive Specification: DiD Coefficient Stability

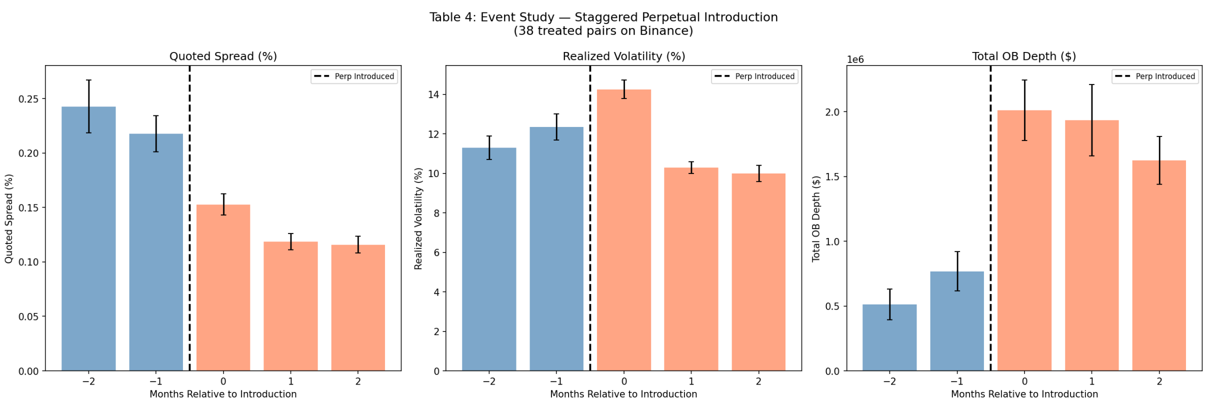

Figure 14: Staggered Introduction: Event Study and Pre/Post Comparison

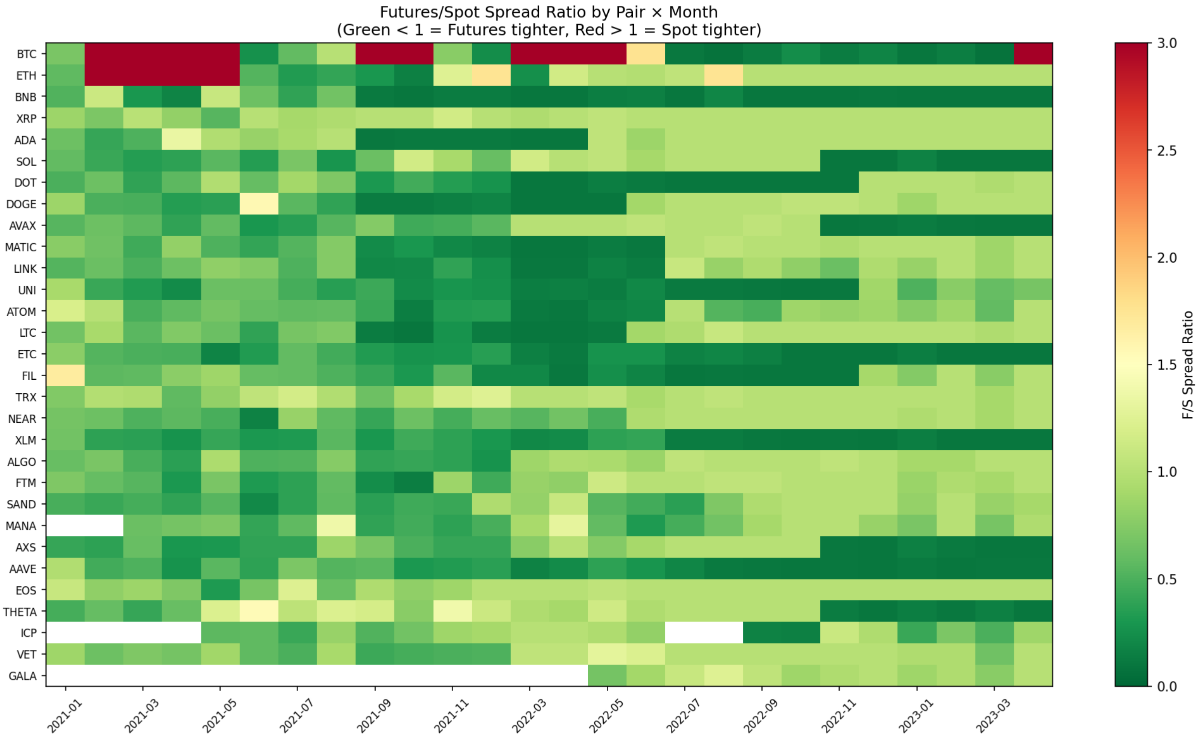

Figure 15: Spot/Perpetual Spread Ratio Heatmap by Pair and Month

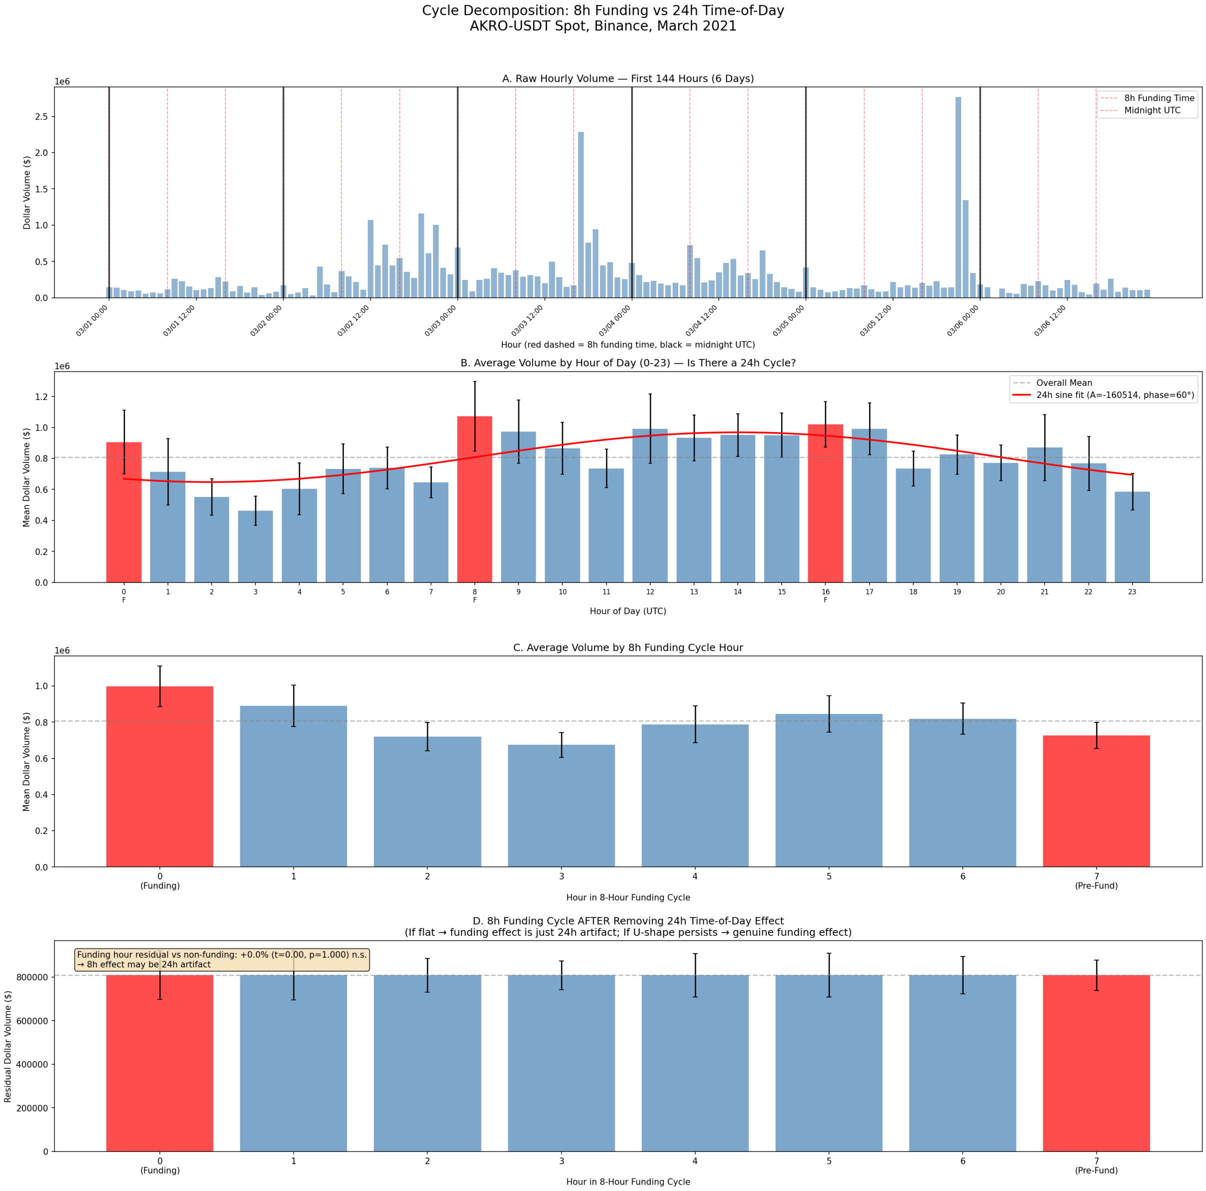

Figure 16: Funding Cycle Decomposition: 8-hour vs. 24-hour Patterns