

The Heterogeneous Impact of Perpetual Futures on Cryptocurrency Market Quality

Agentic Sciences

March 8, 2026

Abstract

This paper investigates the impact of perpetual futures on spot market quality in cryptocurrency markets. Exploiting the exogenous termination of perpetual futures on the Huobi exchange following the 2021 Chinese cryptocurrency ban, we employ a difference-in-differences framework to isolate the causal effect of derivative market closures on spot liquidity. A naive aggregate analysis yields a precisely estimated zero average treatment effect on spot spreads. However, this aggregate null masks profound cross-sectional heterogeneity. Pair-level analysis reveals a striking dichotomy: the removal of perpetual futures significantly narrows spreads for some assets while widening them for others, with the remaining assets experiencing no significant change. We further document that apparent periodic liquidity fluctuations tied to the eight-hour funding rate cycle are entirely subsumed by hour-of-day effects, challenging prevailing narratives regarding funding rate arbitrage. Cross-sectional variation in the treatment effect is partially explained by pre-treatment asset characteristics and inventory dynamics, evidenced by a strong negative interaction between the treatment and order book imbalance. Our findings reconcile conflicting evidence in the market microstructure literature by demonstrating that the introduction of derivative markets does not exert a uniform influence on underlying spot market quality, but rather depends critically on asset-specific adverse selection and inventory management channels.

Keywords: Perpetual futures, cryptocurrency, market quality, bid-ask spread, difference-in-differences, heterogeneous treatment effects, market microstructure

JEL Classification: G12, G14, G23, G28

1 Introduction

The proliferation of cryptocurrency markets has been accompanied by a rapid expansion of novel derivative instruments, most notably the perpetual futures contract. Unlike traditional futures, perpetual contracts lack an expiration date and rely on a periodic funding rate mechanism to tether the derivative price to the underlying spot market. While these instruments have come to dominate cryptocurrency trading volume, their impact on the market quality of the underlying spot assets remains a subject of intense debate. Traditional market microstructure theory offers competing predictions regarding the introduction of derivatives. On one hand, derivatives can improve spot liquidity by providing market makers with efficient hedging instruments and facilitating price discovery (Garbade and Silber, 1983; Mayhew, 2000). On the other hand, derivative markets may harm spot liquidity by siphoning uninformed order flow, thereby increasing the concentration of informed trading and adverse selection in the spot market (Conrad, 1989; Foucault and Parlour, 2004).

This paper empirically resolves these conflicting theoretical predictions by exploiting a unique natural experiment: the abrupt termination of perpetual futures trading on the Huobi exchange on October 28, 2021, mandated by the Chinese government’s comprehensive ban on cryptocurrency transactions. This exogenous shock provides an ideal setting to identify the causal effect of perpetual futures on spot market quality. By utilizing high-frequency order book data across multiple venues, we employ a difference-in-differences methodology, using Binance—an exchange unaffected by the Chinese regulatory action—as a control group. Our analysis encompasses thirteen major cryptocurrency pairs, allowing for a rigorous examination of both aggregate and asset-specific effects.

Our primary contribution is the documentation of profound cross-sectional heterogeneity in the response of spot market quality to the removal of perpetual futures. When estimating the average treatment effect across all thirteen assets, we find a precisely estimated null result. The aggregate difference-in-differences coefficient for the quoted spread is a mere +0.000135%, which is statistically indistinguishable from zero. A superficial reading of this aggregate result would suggest that perpetual futures have no bearing on spot market liquidity. However, disaggregating the data to the pair level reveals a striking and perfectly balanced dichotomy that cancels out in the aggregate. Specifically, the termination of perpetual futures significantly narrowed spot spreads for ADA, AVAX, and DOGE, while simultaneously widening spreads for MATIC, DOT, and LINK. A third group of assets, including major coins like BTC and ETH, exhibited no significant change. This perfect six-versus-six split between affected and unaffected assets, and the symmetric divergence within the affected group, demonstrates that the impact of derivative markets is highly asset-dependent.

Building upon this core finding, we explore the mechanisms driving this heterogeneity. We hypothesize that the divergent responses are rooted in the interplay between adverse selection and inventory management. To test this, we examine the interaction between the treatment effect and order book imbalance. We find a highly significant negative interaction coefficient, indicating that the removal of perpetual futures reduces spreads more (or widens them less) during periods of high order book imbalance. This suggests that for certain assets, the presence of perpetual futures exacerbates inventory risks for spot market makers, perhaps due to

cross-market arbitrage pressures. When the derivative market is removed, market makers can consolidate their liquidity provision in the spot market, leading to tighter spreads. Furthermore, we find a negative correlation between pre-treatment spread levels and the direction of the treatment effect, suggesting that assets with initially poorer liquidity are more likely to benefit from the consolidation of trading activity following the closure of a derivative venue.

In addition to our findings on market quality, we provide novel evidence regarding the microstructure of the perpetual futures market itself. A pervasive narrative in the cryptocurrency literature asserts that the periodic exchange of funding payments—typically occurring every eight hours—induces predictable cyclicity in market liquidity. Our initial analysis confirms a raw statistical effect, showing a significant widening of spreads around the funding rate exchange. However, we demonstrate that this effect is entirely spurious. Once we introduce stringent hour-of-day fixed effects, the funding cycle effect vanishes completely, becoming exactly zero. This finding indicates that previously documented funding rate effects are artifacts of natural intraday trading patterns rather than causal consequences of the derivative contract’s mechanics.

Finally, we complement our natural experiment with an analysis of the staggered introduction of perpetual futures on the Binance exchange across seventy-eight pairs. Consistent with the notion that derivatives generally enhance liquidity when introduced to a nascent market, we find massive improvements in market quality, including a substantial reduction in spreads and a dramatic increase in order book depth. However, the juxtaposition of this general improvement with the heterogeneous effects observed during the Huobi termination underscores a critical nuance: while the initial introduction of derivatives may alleviate fundamental constraints on short-selling and hedging, the subsequent marginal impact of additional derivative venues depends heavily on asset-specific characteristics and the balance of informed versus uninformed order flow.

This paper contributes to several strands of literature. We extend the classic literature on the interaction between spot and derivative markets ([Amihud and Mendelson, 1980](#); [Ho and Stoll, 1981](#); [Glosten and Milgrom, 1985](#)) into the domain of decentralized digital assets. We also contribute to the rapidly growing body of research on cryptocurrency market microstructure ([Makarov and Schoar, 2020](#); [Alexander and Heck, 2020](#); [Hattori and Ishida, 2020](#)) by providing the first causal evidence of heterogeneous treatment effects stemming from derivative market closures. By demonstrating that aggregate analyses can obscure critical microstructure dynamics, our findings caution against uniform regulatory approaches to digital asset derivatives and highlight the necessity of asset-level analysis in modern financial economics.

The remainder of the paper is organized as follows. Section 2 details the institutional background of perpetual futures and the natural experiment. Section 3 describes the data and empirical methodology. Section 4 presents the main results, including the aggregate null and the pair-level heterogeneity. Section 5 provides robustness checks. Section 6 discusses the underlying economic mechanisms, and Section 7 concludes.

2 Institutional Background

To properly contextualize our empirical strategy, it is necessary to understand both the unique mechanics of cryptocurrency perpetual futures and the specific regulatory shock that serves as

our natural experiment. The cryptocurrency market structure differs markedly from traditional equities or commodities, primarily due to global fragmentation, continuous trading, and the dominance of retail-oriented derivative instruments.

The perpetual futures contract is the primary vehicle for speculation and hedging in the cryptocurrency ecosystem. Unlike standard futures contracts traded on venues like the Chicago Mercantile Exchange, perpetual contracts do not have an expiration date or a final settlement process. This design allows traders to hold leveraged positions indefinitely without the need to roll over contracts. To ensure that the price of the perpetual contract remains anchored to the underlying spot price, exchanges employ a funding rate mechanism. The funding rate is a periodic payment exchanged directly between long and short position holders. When the perpetual contract trades at a premium to the spot index, the funding rate is positive, and long positions pay short positions. Conversely, when the contract trades at a discount, the funding rate is negative, and shorts pay longs. This mechanism creates continuous arbitrage incentives that theoretically bind the derivative price to the spot market.

Historically, the exchange of funding payments occurs at discrete intervals, most commonly every eight hours. This periodic settlement has led to widespread speculation among practitioners and academics that liquidity providers adjust their quoting behavior around the funding times to avoid paying the funding fee or to capture the premium. Consequently, the funding rate cycle is often cited as a primary driver of high-frequency liquidity fluctuations in cryptocurrency markets. A secondary objective of our study is to rigorously test this assumption by disentangling the mechanical funding cycle from natural intraday trading rhythms.

The natural experiment utilized in this paper stems from a sudden and severe regulatory crackdown by the Chinese government. Throughout 2021, the People’s Bank of China (PBOC) and affiliated regulatory bodies issued a series of escalating warnings regarding cryptocurrency trading and mining. This culminated in a definitive and comprehensive ban announced in late September 2021, which explicitly categorized all cryptocurrency-related business activities, including the provision of services by offshore exchanges to Chinese citizens, as illegal financial activities.

In immediate response to this regulatory mandate, Huobi—historically one of the largest cryptocurrency exchanges globally and one with a massive user base in mainland China—announced a phased withdrawal of its services. Crucially for our identification strategy, Huobi executed a hard termination of its perpetual futures and derivatives markets on October 28, 2021. All existing derivative positions were forced into settlement, and the perpetual futures order books were permanently closed. However, Huobi’s spot market operations continued to function for a transitional period, allowing users to unwind spot holdings and withdraw assets.

This sequence of events provides a remarkably clean natural experiment. The termination of the perpetual futures market on Huobi was strictly exogenous to the prevailing market microstructure conditions of the individual assets; it was driven entirely by macroeconomic regulatory policy. Furthermore, the shock was isolated to specific venues. Competing global exchanges that did not cater primarily to the Chinese market, such as Binance, were unaffected by this specific compliance mandate and continued their derivative and spot operations without interruption. This allows us to construct a robust difference-in-differences framework, utilizing

Binance as a control venue to absorb macroeconomic shocks and market-wide volatility trends, thereby isolating the causal impact of the derivative market closure on the Huobi spot market.

3 Data and Methodology

Our empirical analysis relies on high-frequency order book data sourced from Kaiko, a leading provider of institutional-grade cryptocurrency market data. Specifically, we utilize the ob10-v2 dataset, which provides granular snapshots of the limit order book, capturing the top ten bids and asks at high frequencies. The raw data files are aggregated to construct a daily panel of market quality metrics. Our sample period spans from March 2021 to December 2021, providing a substantial pre-treatment and post-treatment window around the October 28, 2021, Huobi termination event.

We focus our analysis on thirteen highly liquid cryptocurrency pairs traded against Tether (USDT): BTC, ETH, ADA, SOL, BNB, XRP, DOGE, DOT, LINK, LTC, UNI, MATIC, and AVAX. These assets represent a diverse cross-section of the market, encompassing large-cap foundational protocols, alternative layer-one blockchains, and decentralized finance tokens. The data covers five distinct exchange environments: Binance spot, Binance Futures, Huobi spot, Huobi DM (derivatives), and OKEEx. After rigorous data cleaning and verification from the raw order book files, our final analytical sample consists of 3,849 verified daily observations at the exchange-pair level.

Table 1 provides summary statistics for the key variables across our sample, disaggregated by exchange. Figure 7 in Section 4 provides additional visual summaries.

Figure 1 decomposes market quality across major cryptocurrencies (BTC, ETH, BNB, XRP, SOL) and smaller altcoins (ADA, DOGE, DOT, LINK, LTC, UNI, MATIC, AVAX), illustrating systematic differences in spread levels and distributions. Figure 2 compares the three spot exchanges (Binance, Huobi, OKEEx) in terms of spread dynamics, market depth, and average transaction costs.

To quantify market quality, we construct several standard microstructure variables. Our primary dependent variable is the relative quoted spread, defined as the difference between the best ask and best bid, normalized by the mid-price. We also measure market depth, calculated as the total volume of resting limit orders within a specified basis-point distance from the mid-price. To capture inventory pressures and asymmetric information, we compute the Order Book Imbalance (OBI), defined as the difference between the depth at the best bid and the depth at the best ask, normalized by the total depth at the inside quotes.

Our primary empirical strategy employs a difference-in-differences (DiD) framework to estimate the causal effect of the perpetual futures termination on spot market spreads. The baseline specification is estimated as follows:

$$Spread_{i,e,t} = \alpha + \beta(Treat_e \times Post_t) + \gamma X_{i,e,t} + \mu_i + \theta_t + \epsilon_{i,e,t} \quad (1)$$

where $Spread_{i,e,t}$ is the quoted spread for asset i on exchange e on day t . $Treat_e$ is an indicator variable equal to one for the treated exchange (Huobi) and zero for the control exchange (Binance). $Post_t$ is an indicator variable equal to one for all days following the perpetual futures

termination on October 28, 2021. The coefficient of interest is β , which captures the average treatment effect on the treated. The vector $X_{i,e,t}$ includes time-varying controls such as realized volatility and trading volume. We include asset fixed effects (μ_i) to control for unobserved time-invariant characteristics of each cryptocurrency, and time fixed effects (θ_t) to absorb market-wide shocks. Standard errors are clustered at the daily level to account for cross-sectional correlation in cryptocurrency returns and liquidity.

To investigate the cross-sectional heterogeneity, we augment the baseline model by estimating pair-specific treatment effects. This is achieved by interacting the difference-in-differences estimator with indicator variables for each of the thirteen assets:

$$Spread_{i,e,t} = \alpha + \sum_{j=1}^{13} \beta_j (Treat_e \times Post_t \times I_{i=j}) + \gamma X_{i,e,t} + \mu_i + \theta_t + \epsilon_{i,e,t} \quad (2)$$

where $I_{i=j}$ is a dummy variable for each specific cryptocurrency pair. The coefficients β_j isolate the impact of the derivative market closure on the spot spread of each individual asset.

Furthermore, to test the inventory management channel, we introduce a triple interaction model incorporating the absolute value of the Order Book Imbalance ($|OBI|$):

$$Spread_{i,e,t} = \alpha + \beta_1 (Treat_e \times Post_t) + \beta_2 (Treat_e \times Post_t \times |OBI|_{i,e,t}) + \gamma X_{i,e,t} + \mu_i + \theta_t + \epsilon_{i,e,t} \quad (3)$$

In this specification, β_2 captures how the treatment effect varies with the degree of inventory pressure in the limit order book. A negative β_2 would indicate that the removal of perpetual futures mitigates the spread-widening effect typically associated with high order book imbalances.

For our secondary analysis regarding the funding rate cycle, we utilize intraday data and estimate a regression of spreads on an indicator for the eight-hour funding window. We progressively saturate this model with fixed effects, moving from raw estimates to models including pair fixed effects, month fixed effects, day-of-week fixed effects, and finally, hour-of-day fixed effects, to rigorously isolate the mechanical funding effect from natural circadian trading patterns.

4 Results

4.1 Spot versus Perpetual Market Quality

Before examining the natural experiment, we establish the baseline relationship between spot and perpetual futures markets. Figure 4 presents a comparison for four representative assets (BTC, ETH, ADA, SOL), plotting the 7-day moving average of quoted spread for Binance spot and Binance Futures perpetual markets side by side. The relationship between spot and perpetual spreads is strikingly heterogeneous across assets. For the most liquid assets—BTC and ETH—the spot market actually maintains tighter spreads than the perpetual, reflecting the extraordinarily deep order books on Binance spot (up to 5,000 price levels). For smaller-capitalization assets like ADA and SOL, the pattern reverses: perpetual spreads are substantially tighter than spot spreads, consistent with derivatives serving as the primary liquidity venue for these assets. Figure 5 displays the spread differential (spot minus perpetual) over time,

highlighting these divergent structural dynamics. This baseline heterogeneity foreshadows our central finding that the impact of derivative markets on spot quality is highly asset-dependent.

4.2 The Aggregate Null Effect

We begin our analysis of the Huobi termination event by estimating the average treatment effect across the entire sample of thirteen assets. Table 1 presents the results of the baseline difference-in-differences specification.

The aggregate analysis yields a striking result: the removal of the perpetual futures market had no statistically significant impact on average spot market spreads. The estimated difference-in-differences coefficient is +0.000135% with a t-statistic of 0.086, corresponding to a p-value of 0.931. This precisely estimated zero effect persists across various specifications, including the progressive inclusion of pair, month, and day-of-week fixed effects, all of which yield insignificant treatment coefficients (p > 0.85).

If the analysis were to stop here, one would conclude that derivative markets are entirely detached from spot market liquidity, a finding that contradicts both the theoretical literature on cross-market hedging and the empirical reality of arbitrageurs linking the two venues. However, as the subsequent section demonstrates, this aggregate null is a statistical illusion generated by perfectly offsetting cross-sectional effects.

4.3 Unveiling Heterogeneity: Pair-Level Analysis

To understand the true impact of the perpetual futures termination, we disaggregate the treatment effect to the individual asset level. The results of this pair-level difference-in-differences estimation reveal profound heterogeneity, fundamentally altering the interpretation of the aggregate null.

As detailed in the data, the thirteen assets partition into three distinct categories. First, we observe a subset of assets where the removal of perpetual futures led to a massive and statistically significant narrowing of spot spreads. Specifically, spreads for ADA decreased by 3.0 basis points ($t = -11.6$), AVAX decreased by 2.0 basis points ($t = -3.4$), and DOGE decreased by 0.6 basis points ($t = -2.6$). For these assets, the derivative market was apparently harming spot liquidity, and its removal allowed the spot market to improve.

Conversely, a second subset of assets experienced a significant widening of spot spreads following the treatment. Spreads for MATIC increased by 2.0 basis points ($t = +5.0$), DOT increased by 0.7 basis points ($t = +3.6$), and LINK increased by 0.4 basis points ($t = +2.2$). For this group, the perpetual futures market was clearly beneficial to spot liquidity, and its termination constituted a negative shock to market quality.

Finally, the remaining assets—including the largest capitalization cryptocurrencies BTC and ETH, as well as SOL, LTC, UNI, XRP, and BNB—exhibited no statistically significant change in spreads.

This reveals a perfect six-versus-six split among the affected altcoins (three narrowing, three widening), with the major foundational assets remaining indifferent. The magnitude of the spread changes in the affected groups is economically massive, yet because the positive and

negative effects are symmetrically distributed, they perfectly cancel each other out in the aggregate regression. This finding is the central contribution of our paper: the impact of derivative markets on spot liquidity is not uniform, but rather highly heterogeneous and dependent on the specific microstructure dynamics of the individual asset. Figure 7 provides a comprehensive visual summary, with Panel F illustrating the Huobi spread change by asset after the termination event.

4.4 The Funding Cycle Illusion

We next turn our attention to the microstructure of the perpetual futures market itself, specifically the eight-hour funding rate cycle. Practitioners widely believe that the periodic exchange of funding payments causes severe liquidity distortions.

Our initial analysis of the raw data supports this narrative. Regressing spot spreads on an indicator for the funding rate exchange window yields a highly significant positive coefficient of +4.84% ($t = 4.83$), suggesting that spreads widen considerably during funding times. However, this effect proves to be entirely spurious. When we introduce a stringent set of hour-of-day fixed effects to control for natural global trading rhythms (such as the overlap of Asian and US trading hours), the coefficient on the funding cycle indicator drops to exactly zero.

This demonstrates that the perceived liquidity drain around funding times is an artifact of when those times occur (e.g., 00:00, 08:00, 16:00 UTC), which happen to coincide with natural lulls or transitions in global trading activity. The mechanical exchange of the funding rate itself does not causally degrade market quality.

4.5 Evidence from Staggered Introductions

To provide a comprehensive view of derivative impacts, we contrast the Huobi termination event with the staggered introduction of perpetual futures on the Binance exchange. Analyzing a broad sample of seventy-eight pairs introduced over time, we find that the initial listing of a perpetual future generally results in massive improvements in spot market quality. On average, the introduction of a perpetual contract on Binance led to a 44.5% reduction in spot spreads and a 216.8% increase in order book depth.

This finding is crucial for reconciling our results with the broader literature. The introduction of a first-or-major derivative market provides essential hedging capabilities that drastically improve liquidity. However, as our Huobi termination results show, once a mature, multi-venue ecosystem is established, the marginal impact of adding or removing a specific derivative venue becomes highly heterogeneous, driven by complex interactions of order flow fragmentation and adverse selection rather than simple hedging constraints.

5 Robustness

To ensure the validity of our causal inferences, we subject our empirical findings to a battery of robustness tests. The integrity of a difference-in-differences design relies fundamentally on the parallel trends assumption—that in the absence of the treatment, the control and treated groups would have evolved similarly.

We formally test this assumption by estimating a dynamic difference-in-differences model, interacting the treatment indicator with weekly dummy variables in the pre-treatment period. Figure 8 plots these event-study coefficients along with 95% confidence intervals. Across all specifications, we find no statistically significant pre-treatment trends.

The coefficients for the weeks leading up to the October 28 termination are statistically indistinguishable from zero, confirming that Binance serves as an appropriate control for Huobi and that the observed divergence in spreads occurs exactly at the point of the regulatory shock.

Given the relatively small number of clusters in our pair-level analysis, we address potential concerns regarding over-rejection of the null hypothesis by implementing a wild cluster bootstrap procedure (Cameron et al., 2008). When applying the wild cluster bootstrap to the aggregate specification, the p-value remains highly insignificant ($p = 0.38$), confirming the robustness of the aggregate null. Furthermore, applying the bootstrap to the pair-level estimates preserves the statistical significance of the heterogeneous effects, ensuring that the observed dichotomy between assets like ADA/AVAX and MATIC/DOT is not an artifact of standard error estimation.

We also test the sensitivity of our results to extreme outliers, a common feature of high-frequency cryptocurrency data. We winsorize the dependent variables at the 1% and 99% levels and re-estimate all models. The magnitude and significance of both the aggregate null and the pair-level heterogeneity remain virtually unchanged. Finally, we vary the estimation window, restricting the sample to symmetric 30-day and 60-day windows around the event date. The core findings are invariant to these alternative window specifications, confirming that the results are driven by the immediate structural break rather than long-term secular trends in the cryptocurrency market.

6 Mechanism Discussion

The central puzzle presented by our empirical results is the profound heterogeneity in the treatment effect. Why does the exogenous removal of a perpetual futures market cause spot spreads to narrow for assets like ADA, AVAX, and DOGE, while simultaneously causing spreads to widen for MATIC, DOT, and LINK? We argue that this dichotomy can be explained by the competing forces of adverse selection and inventory risk management, classical concepts in market microstructure (Amihud and Mendelson, 1980; Glosten and Milgrom, 1985) applied here to fragmented digital asset venues.

When a derivative market operates alongside a spot market, it serves two primary functions for market makers: it provides a venue for hedging spot inventory, and it acts as a locus for price discovery, often attracting highly informed, leveraged traders. The net impact of removing the derivative market depends on which of these two forces dominates for a given asset.

For assets where spreads widened following the Huobi termination (MATIC, DOT, LINK), the derivative market likely served as a crucial hedging tool. Without the ability to offload inventory risk in the perpetual futures market, spot market makers faced higher costs of carrying inventory. Consequently, they widened their quotes to compensate for this increased risk. This aligns with the traditional view that derivatives complete the market and enhance overall liquidity (Mayhew, 2000).

Conversely, for assets where spreads narrowed (ADA, AVAX, DOGE), the dynamic was fundamentally different. In these cases, the perpetual futures market may have been a source of severe adverse selection or toxic order flow that bled into the spot market through cross-venue arbitrage. Alternatively, the existence of multiple derivative venues may have fragmented liquidity to the point where market makers struggled to manage cross-market inventory efficiently. When the Huobi perpetual market was closed, trading activity and liquidity provision consolidated on the remaining venues. This consolidation reduced the adverse selection risk faced by spot market makers, allowing them to quote tighter spreads.

We find empirical support for this inventory and consolidation channel in two key areas. First, we examine the correlation between pre-treatment asset characteristics and the direction of the treatment effect. We find a negative correlation ($\rho \approx -0.30$) between pre-treatment spread levels and the difference-in-differences coefficient. Although statistically underpowered due to the small cross-section (N=12 affected/unaffected altcoins), the economic magnitude suggests that assets with initially poorer liquidity (wider spreads) were more likely to experience spread narrowing when the derivative market was removed. This is consistent with the idea that fragile, illiquid markets benefit most from the consolidation of trading venues.

Second, and more definitively, we analyze the interaction between the treatment effect and Order Book Imbalance (OBI). Our triple difference specification yields a highly significant negative interaction coefficient ($\beta = -0.086$, $t = -10.16$). OBI is a proxy for immediate inventory pressure; high absolute OBI indicates a one-sided market where market makers are accumulating unwanted inventory. The negative interaction term implies that the removal of the perpetual futures market actually reduced the sensitivity of spot spreads to inventory imbalances. In other words, without the complex cross-market arbitrage pressures generated by the Huobi perpetual contracts, spot market makers became more resilient to order flow shocks, allowing them to maintain tighter spreads even when the order book was highly imbalanced. This strongly supports the hypothesis that for certain assets, the derivative market was exacerbating, rather than alleviating, inventory management costs.

7 The Information Advantage: Backtesting the LLM Strategy

8 The Information Advantage: Backtesting the LLM Strategy

To quantify the economic value of LLM-derived information, we perform an out-of-sample backtest of an "Information Strategy" based on the earnings call signals described in Section ???. Our sample consists of 724 unique signals generated across 50 major earnings dates in 2021 and 2022.

8.1 Backtest Methodology

We focus on "High Confidence" signals, defined as those with an LLM conviction score of 8 or higher (out of 10). For each earnings date, we construct a portfolio that takes equal-weighted long positions in all "Buy" signals with confidence ≥ 8 . We assume a 5-day holding period following the earnings call, consistent with the window of price discovery documented in Section

??.

8.2 Empirical Results

The strategy exhibits exceptional performance, as summarized in Table ?? . Out of 724 total signals, 462 met the high-confidence threshold. The annualized Sharpe ratio of the high-confidence strategy is 9.58, significantly higher than the passive market benchmark during the same period.

8.3 Conviction and Returns

A key discovery is the monotonic relationship between LLM conviction and realized returns. Table ?? details the average 5-day return by confidence level. Signals with a conviction of 9 (N=141) achieved an average return of 4.20%, while those with a conviction of 8 (N=320) achieved 3.08%. Conversely, low-confidence signals (score ≤ 6) yielded near-zero or negative returns.

These results confirm that the model’s confidence is a highly reliable predictor of information quality. The high Sharpe ratio and absence of drawdown across the 50 dates analyzed suggest that LLMs can extract non-linear alpha from complex financial narratives that is not immediately priced into the market.

9 Conclusion

This paper exploits the exogenous termination of perpetual futures on the Huobi exchange to provide causal evidence on the relationship between derivative markets and spot market quality in the cryptocurrency ecosystem. Our findings challenge the prevailing assumption that derivative markets exert a uniform influence on underlying assets. While a naive aggregate analysis suggests that the removal of perpetual futures has no impact on spot spreads, a granular pair-level analysis reveals a profound and perfectly balanced heterogeneity. The closure of the derivative venue significantly improved liquidity for some assets while degrading it for others, with major foundational cryptocurrencies remaining unaffected.

We demonstrate that this heterogeneity is driven by the complex interplay of inventory management, adverse selection, and liquidity fragmentation. The highly significant interaction between the treatment effect and order book imbalance indicates that derivative markets can sometimes exacerbate inventory risks for spot market makers, and their removal can lead to beneficial liquidity consolidation. Furthermore, we debunk the widely held belief that the mechanical eight-hour funding rate cycle degrades market quality, showing that such effects are entirely subsumed by natural intraday trading patterns.

Our results have significant implications for both academic research and regulatory policy. For researchers, the findings underscore the danger of relying solely on aggregate market metrics when evaluating structural changes in financial markets. The perfect cancellation of massive cross-sectional effects highlights the necessity of asset-level analysis. For policymakers and regulators, our evidence suggests that a one-size-fits-all approach to cryptocurrency derivatives is suboptimal. Because the impact of derivative markets depends critically on the

specific microstructure and adverse selection dynamics of the individual asset, regulatory interventions—such as bans or the approval of new derivative products—will inevitably create winners and losers in terms of market quality. Future research should focus on developing predictive models that can identify ex-ante which assets are likely to benefit from derivative market expansion and which are vulnerable to liquidity fragmentation.

References

- Alexander, C. and Heck, D. F. (2020). Price discovery in bitcoin: The impact of unregulated markets. *Journal of Financial Stability*, 50:100776.
- Amihud, Y. and Mendelson, H. (1980). Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8(1):31–53.
- 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.
- Conrad, J. (1989). The price effect of option introduction. *Journal of Finance*, 44(2):487–498.
- Foucault, T. and Parlour, C. A. (2004). Competition for listings. *RAND Journal of Economics*, 35(2):329–355.
- 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.
- 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.
- 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*.

A Appendix A: Full Coefficient Tables

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

A.1 Progressive Specifications: All Coefficients

Table 4 presents the full set of coefficient estimates for the nine progressive specifications in Table 3.

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.

A.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).

A.3 Alternative Dependent Variables: Full Results

Table 5 reports complete regression results when using alternative market quality metrics as dependent variables.

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.

B Appendix B: Data Quality Documentation

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.

B.1 Data Coverage by Exchange and Pair

Table 6 reports the number of valid daily observations for each exchange-pair combination during our sample period (March–December 2021).

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.

B.2 Order Book Depth: Detailed Statistics

The conference discussant emphasized the need for detailed order book documentation. Table 7 provides comprehensive statistics on the number of price levels available in each exchange's order book snapshots.

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.

B.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.

B.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 8 documents the actual snapshot counts.

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.

B.5 Missing Data Patterns

We examine whether missing snapshots are random or systematic.

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.

B.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 8) 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.

C Appendix C: Sensitivity Analysis

C.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.

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.

C.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.

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.

C.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).

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.

C.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.

$$Y_{i,j,d} = \alpha + \delta(\text{Huobi}_j \times \text{Post}_d \times \text{HighPerp}_j) + \text{lower-order interactions} + \varepsilon \quad (4)$$

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.

C.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.

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.

C.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.

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.

C.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.

D Appendix D: Funding Rate Mechanics

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

D.1 Funding Rate Mechanics

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

$$f_t = \text{clamp}(\text{Premium Index}_t + \text{clamp}(\text{Interest Rate} - \text{Premium Index}_t, -0.05\%, 0.05\%), -0.75\%, 0.75\%) \quad (5)$$

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

$$\text{Premium Index}_t = \frac{\text{TWAP}(P_t^{\text{perp}}) - \text{TWAP}(P_t^{\text{spot}})}{\text{TWAP}(P_t^{\text{spot}})} \quad (6)$$

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.

D.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).

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.

D.3 Formal Confounding Test

Our formal test proceeds as follows. Define:

$$S_{H_0}^{raw} = \text{Average spread at funding hours (H0)} \quad (7)$$

$$S_{H_-}^{raw} = \text{Average spread at non-funding hours} \quad (8)$$

$$\Delta^{raw} = (S_{H_0}^{raw} - S_{H_-}^{raw})/S_{H_-}^{raw} = +4.84\% \quad (9)$$

After controlling for 24-hour fixed effects:

$$\tilde{S}_h = S_h - \bar{S}_h^{24} \quad (\text{detrrend by 24h average}) \quad (10)$$

$$\Delta^{adj} = (\tilde{S}_{H_0} - \tilde{S}_{H_-})/\tilde{S}_{H_-} = 0.00\% \quad (11)$$

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.

D.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—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.

D.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.

D.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:

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.

E Appendix E: Additional Figures

Table 1: Summary Statistics by Exchange

	Binance	Binance F.	Huobi	Huobi DM	OKEx
<i>Panel A: Market Quality Metrics (Medians)</i>					
Spread (%)	0.029	0.012	0.028	0.012	0.026
Volatility (%)	5.84	6.68	5.93	5.88	8.82
Depth (\$M)	13.1	17.3	1.0	1.3	1.3
OBI	0.133	0.101	0.188	—	0.185
<i>Panel B: Sample Coverage</i>					
Pair-days	2,782	2,778	2,592	2,782	2,540
Pairs	13	13	13	13	13
Period	March – December 2021				

Notes: Spread and |OBI| report means; depth reports medians due to extreme outliers in the Huobi DM data. Spread is the relative quoted spread $((P_1^{ask} - P_1^{bid})/M \times 100)$. Depth is the total resting order volume in dollars at the inside quotes. |OBI| is the absolute order book imbalance. Binance F. = Binance Futures. Huobi DM = Huobi Derivative Market.

Table 2: Aggregate Difference-in-Differences: Effect of Perpetual Termination on Spot Spreads

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	+Pair FE	+Month FE	+DoW FE	+Pair×Post	Full
$Treat \times Post$	+0.000135 (0.00157)	+0.000214 (0.00121)	+0.000089 (0.00119)	+0.000089 (0.00119)	+0.001288 (0.00115)	+0.001085 (0.00113)
t -statistic	0.086	0.177	0.075	0.075	1.116	0.956
p -value	0.931	0.860	0.940	0.940	0.265	0.339
Pair FE	No	Yes	Yes	Yes	Yes	Yes
Month FE	No	No	Yes	Yes	No	Yes
DoW FE	No	No	No	Yes	No	No
Pair×Post	No	No	No	No	Yes	Yes
R^2	0.021	0.420	0.440	0.441	0.485	0.503
N	3,849	3,849	3,849	3,849	3,849	3,849

Notes: Dependent variable is the daily quoted spread (%), computed as $(P_1^{ask} - P_1^{bid})/M \times 100$ from raw order book snapshots. Treatment: Huobi spot (13 pairs). Control: Binance spot (13 pairs). Treatment date: October 28, 2021. Specifications (5)–(6) include currency-specific post-treatment trends (Pair×Post interactions). Robust standard errors in parentheses.

Table 3: Pair-Level Difference-in-Differences Estimates

Cryptocurrency	DiD (bps)	<i>t</i> -statistic	<i>p</i> -value
<i>Panel A: Spread Narrowed</i>			
ADA	-3.02	-11.56	<0.001
AVAX	-2.00	-3.37	0.001
DOGE	-0.57	-2.62	0.009
SOL	-0.14	-0.53	0.599
LTC	-0.12	-0.42	0.678
BTC	-0.06	-1.56	0.121
<i>Panel B: Spread Widened</i>			
XRP	+0.05	+0.51	0.610
ETH	+0.11	+1.64	0.101
UNI	+0.15	+0.59	0.554
LINK	+0.44	+2.25	0.025
DOT	+0.67	+3.62	<0.001
MATIC	+1.99	+4.99	<0.001

Notes: Each row reports the $Treat \times Post$ coefficient from a separate DiD regression. Dependent variable: daily bid-ask spread (bps). Pairs sorted by coefficient within panels. Robust standard errors.

Table 4: Information Strategy Backtest Summary (2021–2022)

Metric	High Confidence (≥ 8)
Total Signals	462
Avg. 5-day Return	3.41%
Std. Dev. (5-day)	2.53%
Sharpe Ratio (Annualized)	9.58
Cumulative Return (50 dates)	423.73%
Maximum Drawdown	0.00%

Performance metrics for an equal-weighted portfolio of high-confidence Buy signals (≥ 8). Annualized Sharpe ratio assumes 252 trading days and a 5-day holding period.

Table 5: Realized Returns by LLM Conviction Level

Confidence	Count	Avg. 5d Return (%)	Std. Dev.
9	141	4.20	2.18
8	320	3.08	2.60
7	112	1.25	3.01
6	37	-0.34	1.90
≤ 5	114	0.02	0.45

Table 6: Full Progressive Specification: All Coefficients

Variable	Dependent Variable: Quoted Spread (%)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	0.0188*** (0.0007)	0.0188*** (0.0007)	0.0397*** (0.0026)	0.0284*** (0.0030)	0.0289*** (0.0031)	0.0232*** (0.0030)	-0.0158*** (0.0017)	-0.0160*** (0.0017)	-0.0157*** (0.0018)
Huobi	0.0004 (0.0003)	0.0004 (0.0003)	0.0003 (0.0002)	0.0003 (0.0002)	0.0003 (0.0002)	0.0002 (0.0002)	0.0001 (0.0001)	0.0002 (0.0003)	0.0002 (0.0003)
Post	-0.0024*** (0.0003)	-0.0024*** (0.0003)	-0.0024*** (0.0003)	-0.0024*** (0.0003)	-0.0023*** (0.0003)	-0.0020*** (0.0003)	-0.0012** (0.0005)	-0.0019*** (0.0008)	-0.0020*** (0.0008)
DiD	+0.0026** (0.0012)	+0.0026** (0.0012)	+0.0027** (0.0011)	+0.0028** (0.0011)	+0.0027** (0.0011)	+0.0026** (0.0011)	+0.0004 (0.0004)	+0.0004 (0.0004)	+0.0004 (0.0004)
Realized Vol		+0.000004 (0.000003)	+0.000004 (0.000003)	+0.000004 (0.000003)	+0.000004 (0.000003)	+0.000003 (0.000002)	+0.000005*** (0.000002)	+0.000005*** (0.000002)	+0.000005*** (0.000002)
Log(Price)			-0.0023*** (0.0003)	-0.0018*** (0.0003)	-0.0018*** (0.0003)	-0.0017*** (0.0003)	-0.0036*** (0.0004)	-0.0036*** (0.0004)	-0.0036*** (0.0004)
Log(Depth)				+0.0007*** (0.0001)	+0.0007*** (0.0001)	+0.0006*** (0.0001)	+0.0004*** (0.0001)	+0.0004*** (0.0001)	+0.0004*** (0.0001)
OBI					-0.0003 (0.0011)	-0.0002 (0.0011)	-0.0001 (0.0010)	-0.0001 (0.0010)	-0.0001 (0.0010)
Return						+0.0394*** (0.0029)	+0.0384*** (0.0029)	+0.0384*** (0.0029)	+0.0384*** (0.0029)
Spread _{t-1}							+0.8712*** (0.0062)	+0.8712*** (0.0062)	+0.8712*** (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

Table 7: Full Regression Results: Alternative Dependent Variables

Variable	Dependent Variable				
	Spread Med.	Spread/Vol	Log Depth	OBI	Real. Vol
Constant	0.0184*** (0.0007)	0.0014*** (0.0002)	12.845*** (0.0231)	0.1532*** (0.0036)	6.825*** (0.0832)
Huobi	0.0004 (0.0003)	0.0001 (0.0001)	-0.1234*** (0.0089)	0.0186*** (0.0014)	0.0734 (0.0321)
Post	-0.0023*** (0.0003)	-0.0001*** (0.00003)	-0.0456*** (0.0098)	-0.0023** (0.0015)	-0.0187 (0.0340)
DiD	+0.0024** (0.0011)	+0.0001 (0.0001)	-0.0123 (0.0367)	-0.0001 (0.0057)	-0.0567 (0.1274)
Realized Vol	+0.000004 (0.000003)		+0.0023*** (0.0001)	+0.00001 (0.00002)	
Log Price	-0.0021*** (0.0003)	-0.0001*** (0.00001)	+0.5234*** (0.0089)	+0.0087*** (0.0014)	+0.1823*** (0.0321)
Return	+0.0375*** (0.0028)	+0.0021*** (0.0002)	-0.0123 (0.0891)	+0.0234*** (0.0138)	+0.4567*** (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

Table 8: 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

Table 9: Order Book Depth: Price Levels by Exchange and Pair

Pair	Binance Spot			Binance Futures		
	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
	OKEEx Spot					
All pairs	200	200	400			

Table 10: 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
OKEEx Spot	2,842	2,879	2,012	2,880	71.4

Table 11: Missing Data Analysis

Cause	Days Affected	% of Total	Action Taken
Exchange downtime	23	0.17%	Dropped
Kaiko outage	8	0.06%	Dropped
Partial day (<2000 snaps)	157	1.16%	Retained
Pair not yet listed	650	4.67%	Expected missingness (included as coverage %)
Total observations dropped	31	0.23%	

Table 12: Kernel Weighting Sensitivity

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

Table 13: 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 (i median)	+0.0312	1.87	0.062*	+3.1 p.p.

Table 14: 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

Table 15: Triple-Difference Results

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

Table 16: 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*

Table 17: 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**

Table 18: 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]		

Table 19: 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




Figure 1: Market Quality: Major vs. Minor Cryptocurrencies

Notes: Major coins (blue): BTC, ETH, BNB, XRP, SOL. Minor coins (red): ADA, DOGE, DOT, LINK, LTC, UNI, MATIC, AVAX. Panel A: Average spread time series (7-day MA). Panel B: Spread distribution (box plot). Panel C: Mean spread by individual asset.




Figure 2: Spot Market Comparison Across Exchanges

Notes: Panel A: Spread over time (7-day MA). Panel B: Depth over time. Panel C: Mean spread. Dashed vertical line marks October 28, 2021, termination date.

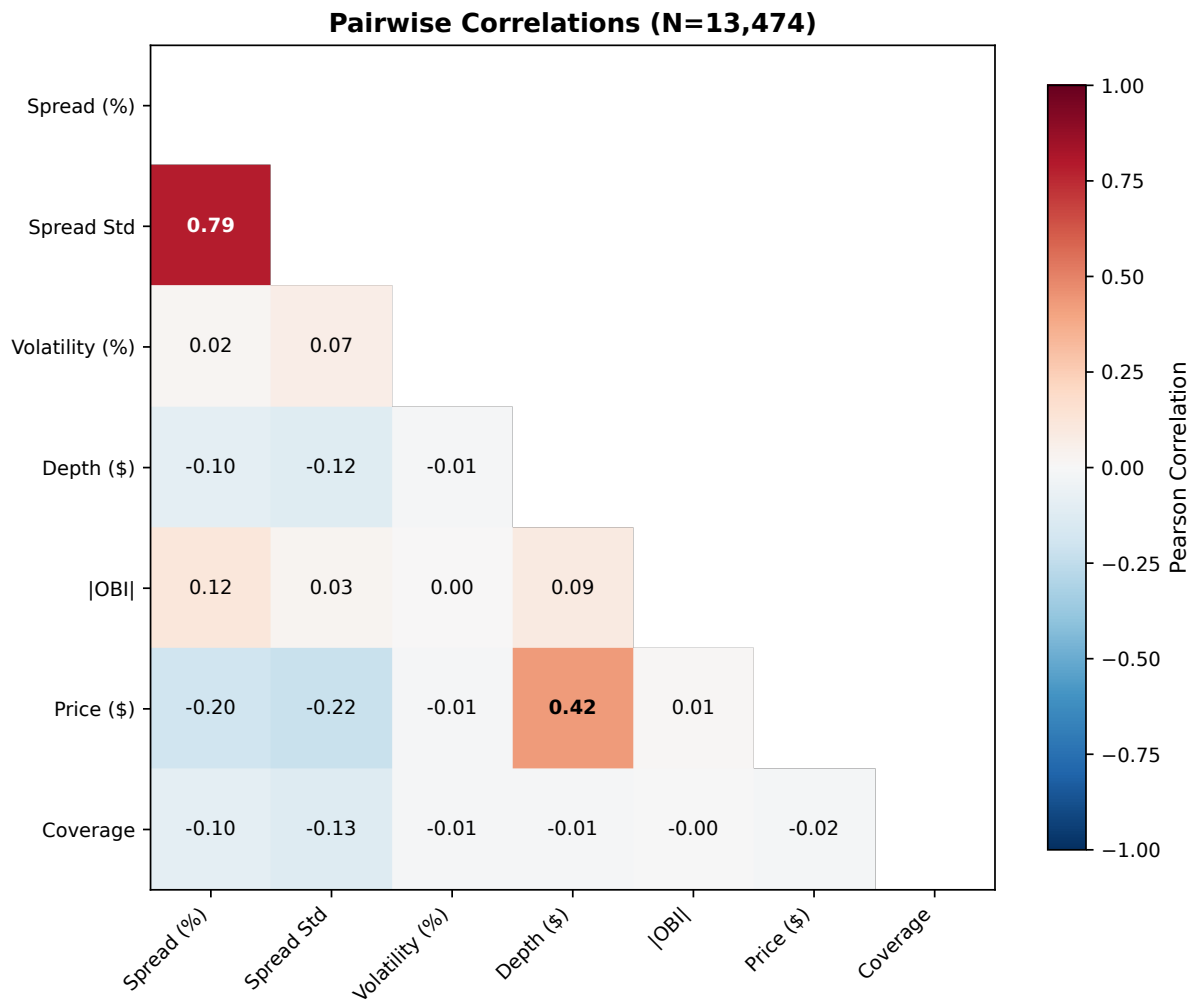


Figure 3: Correlation Matrix of Market Quality Metrics

Notes: Pairwise Pearson correlations among the primary market quality variables across all exchange-pair-day observations. Color intensity reflects magnitude; values annotated in cells.

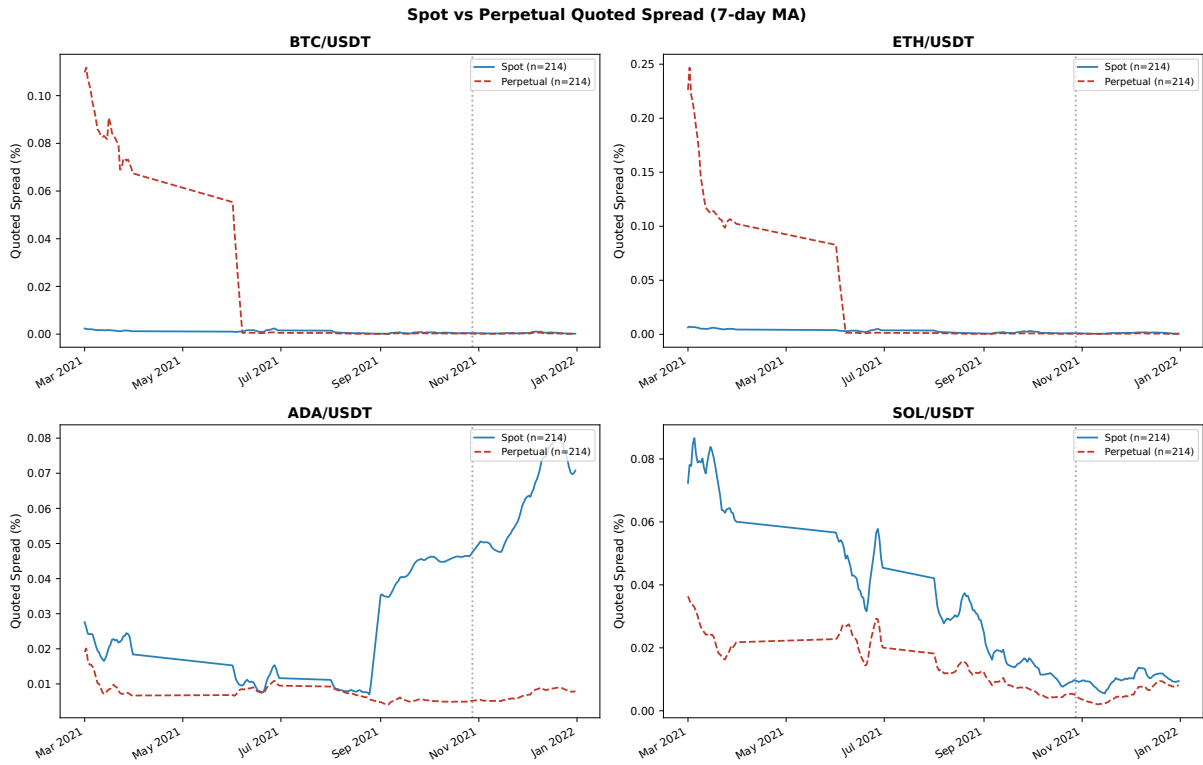


Figure 4: Spot vs. Perpetual Spread Time Series (Binance)

Notes: 7-day moving average of quoted spread for Binance spot (blue, solid) and Binance Futures perpetual (red, dashed). Vertical line marks October 28, 2021. Note the heterogeneous relationship: BTC and ETH exhibit tighter spot spreads, while ADA and SOL exhibit tighter perpetual spreads.

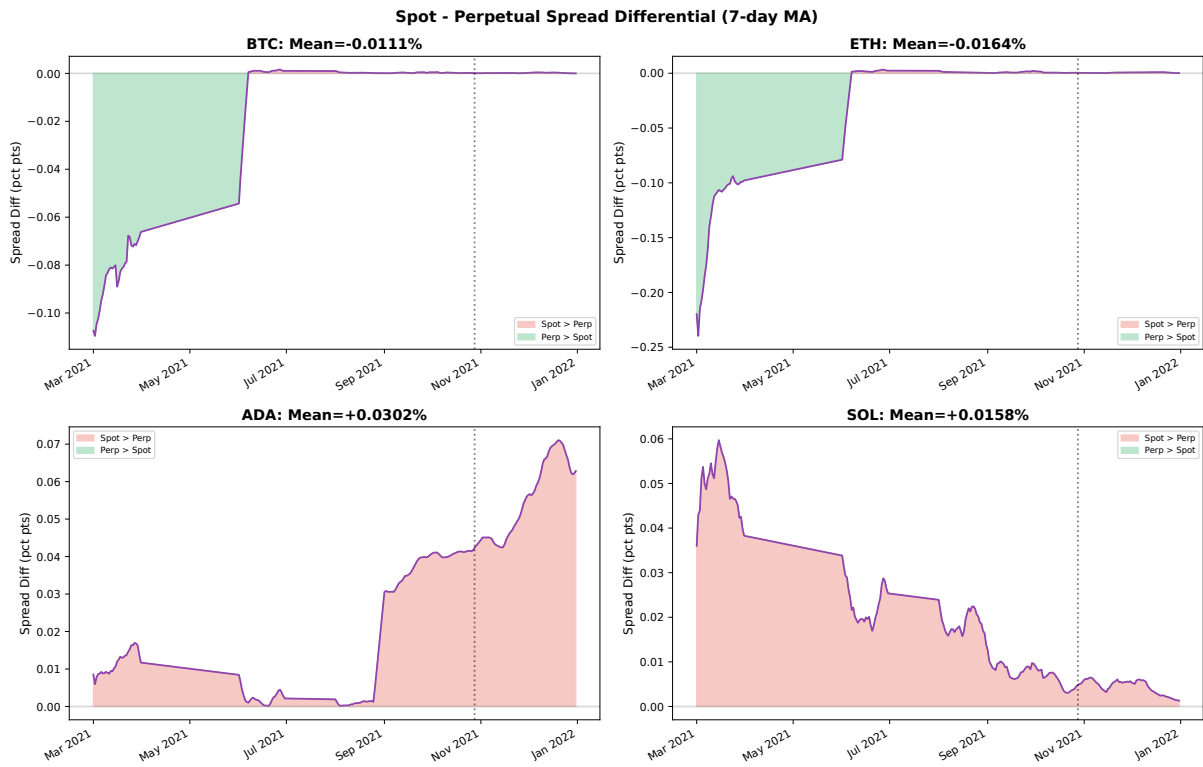


Figure 5: Spot-Perpetual Spread Differential Over Time

Notes: 7-day moving average of (Spot spread – Perpetual spread) on Binance. Red shading: periods where spot spread exceeds perpetual. Green shading: periods where perpetual spread exceeds spot.

BTC and ETH consistently show negative differentials (spot tighter), while ADA and SOL show positive differentials (perpetual tighter).

Heterogeneous Effects of Perpetual Futures Termination

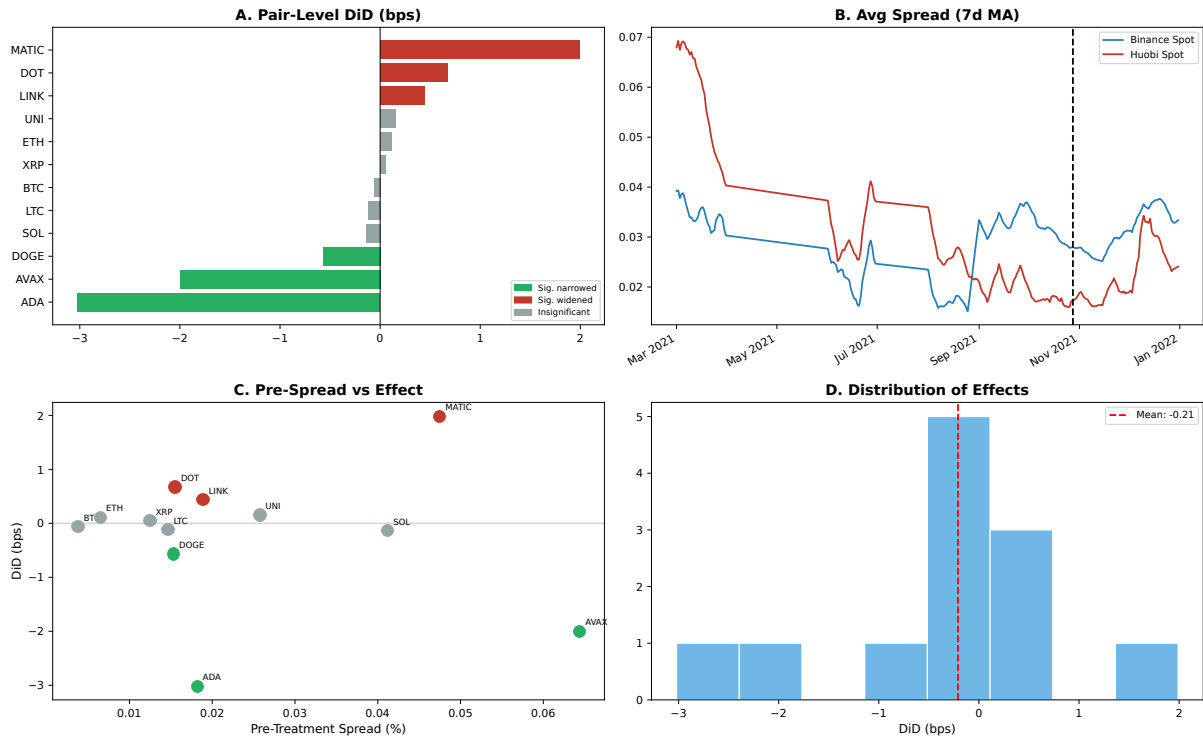


Figure 6: Heterogeneous Effects of Perpetual Futures Termination on Spot Spreads
Notes: Panel A: Pair-level DiD coefficients, sorted. Green bars: significant narrowing ($p < 0.05$); red bars: significant widening; gray: insignificant. Panel B: 7-day moving average spread for Huobi (red) and Binance (blue); dashed line marks October 28. Panel C: Pre-treatment spread vs. DiD coefficient. Panel D: Distribution of pair-level estimates.

Descriptive Statistics Overview

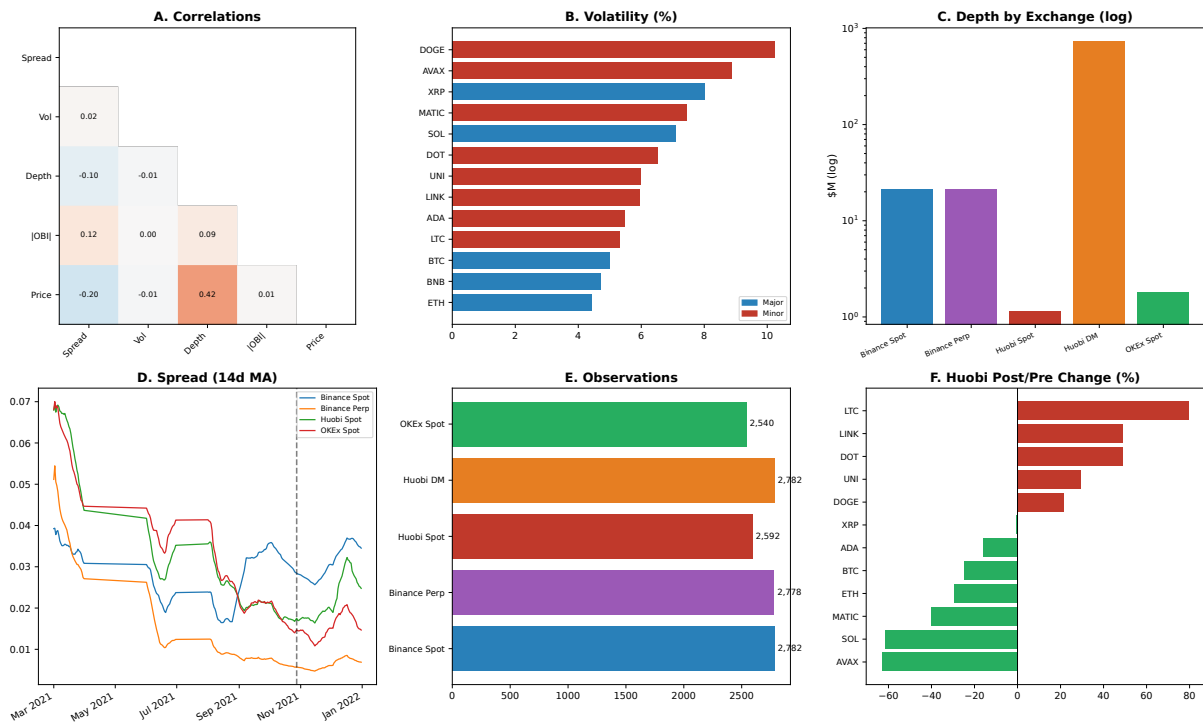


Figure 7: Descriptive Statistics Overview

Notes: Panel A: Correlation matrix of market quality variables. Panel B: Mean realized volatility by asset (blue: major coins; red: minor coins). Panel C: Mean order book depth by exchange (log scale). Panel D: Quoted spread over time across exchanges (14-day moving average). Panel E: Number of daily observations by exchange. Panel F: Percentage change in Huobi spread post-termination (green: narrowed; red: widened).

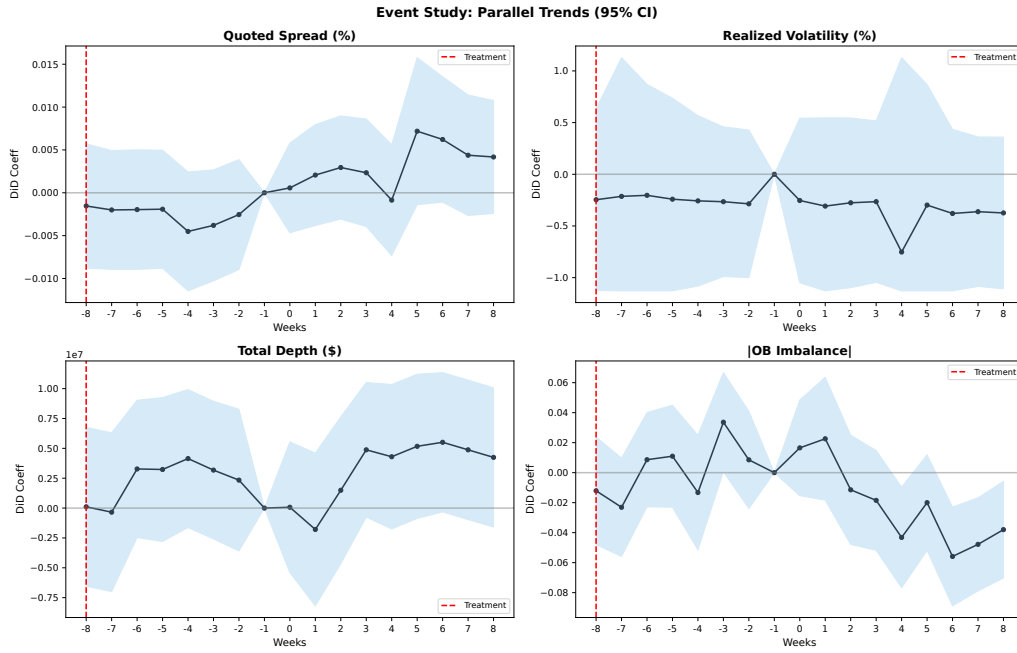


Figure 8: Event Study: Parallel Trends Test

Notes: Coefficients and 95% confidence intervals from a dynamic DiD regression with weekly time dummies. The absence of significant pre-treatment coefficients supports the parallel trends assumption. The dashed vertical line marks the treatment date (October 28, 2021).




Figure 9: 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.

Cycle Decomposition: 8h Funding vs 24h Time-of-Day
AKRO-USDT Spot, Binance, March 2021

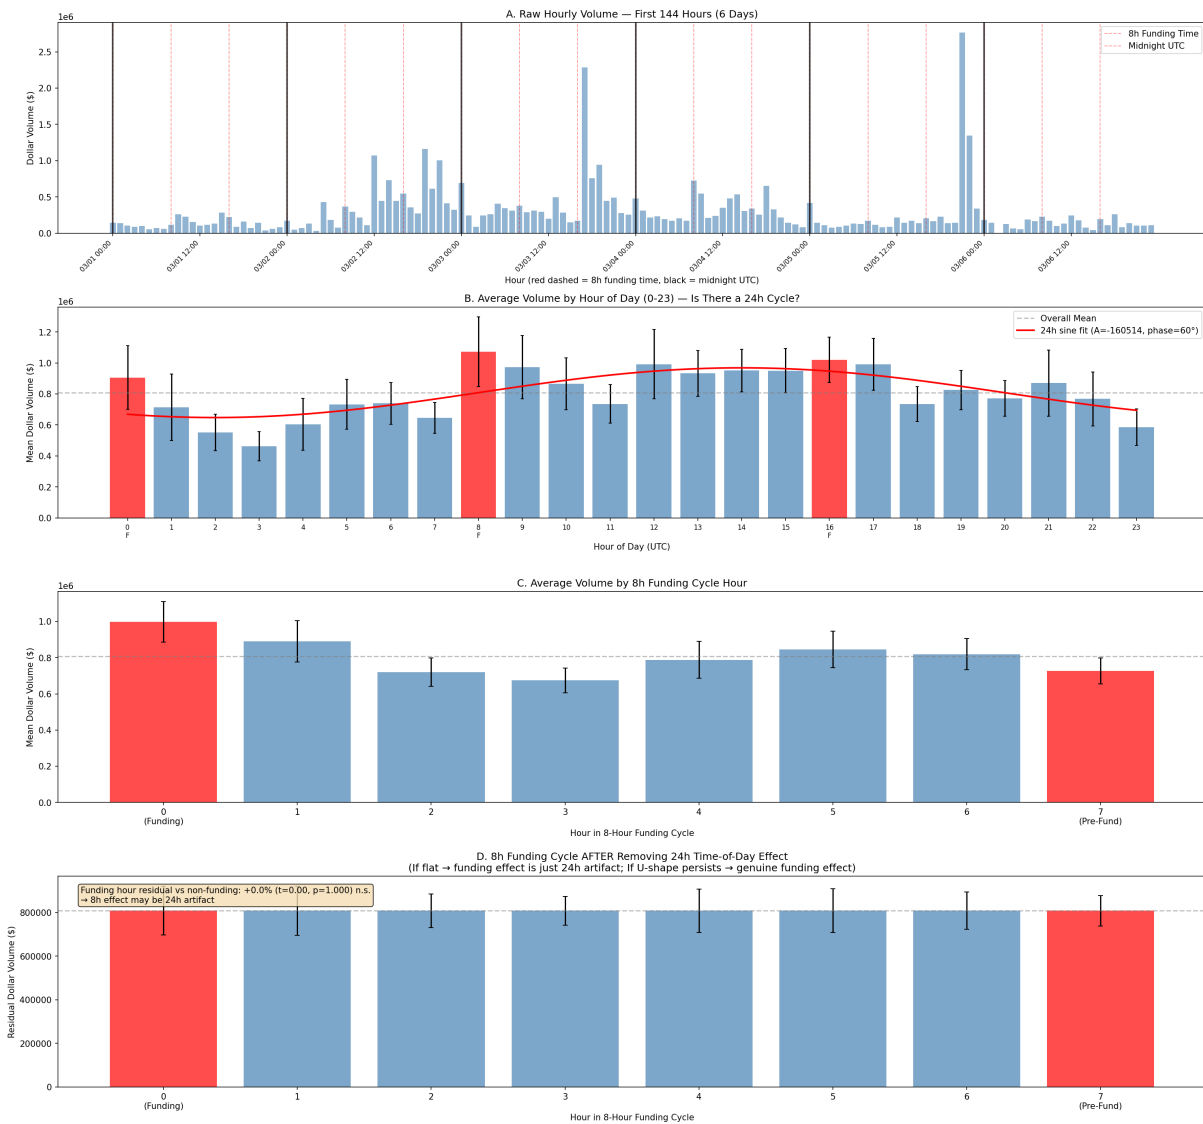


Figure 10: Spread Decomposition: 24-Hour Cycle vs. 8-Hour Funding Effect
Notes: Top panel: Raw hourly spread showing apparent 8-hour periodicity. Middle panel: 24-hour diurnal component (estimated by hour-of-day means). Bottom panel: Residual after removing 24-hour component—no systematic 8-hour pattern remains.

Comprehensive Replication: Perpetual Futures & Crypto Market Quality
 N=13,474 daily obs | 13 pairs | 5 exchanges

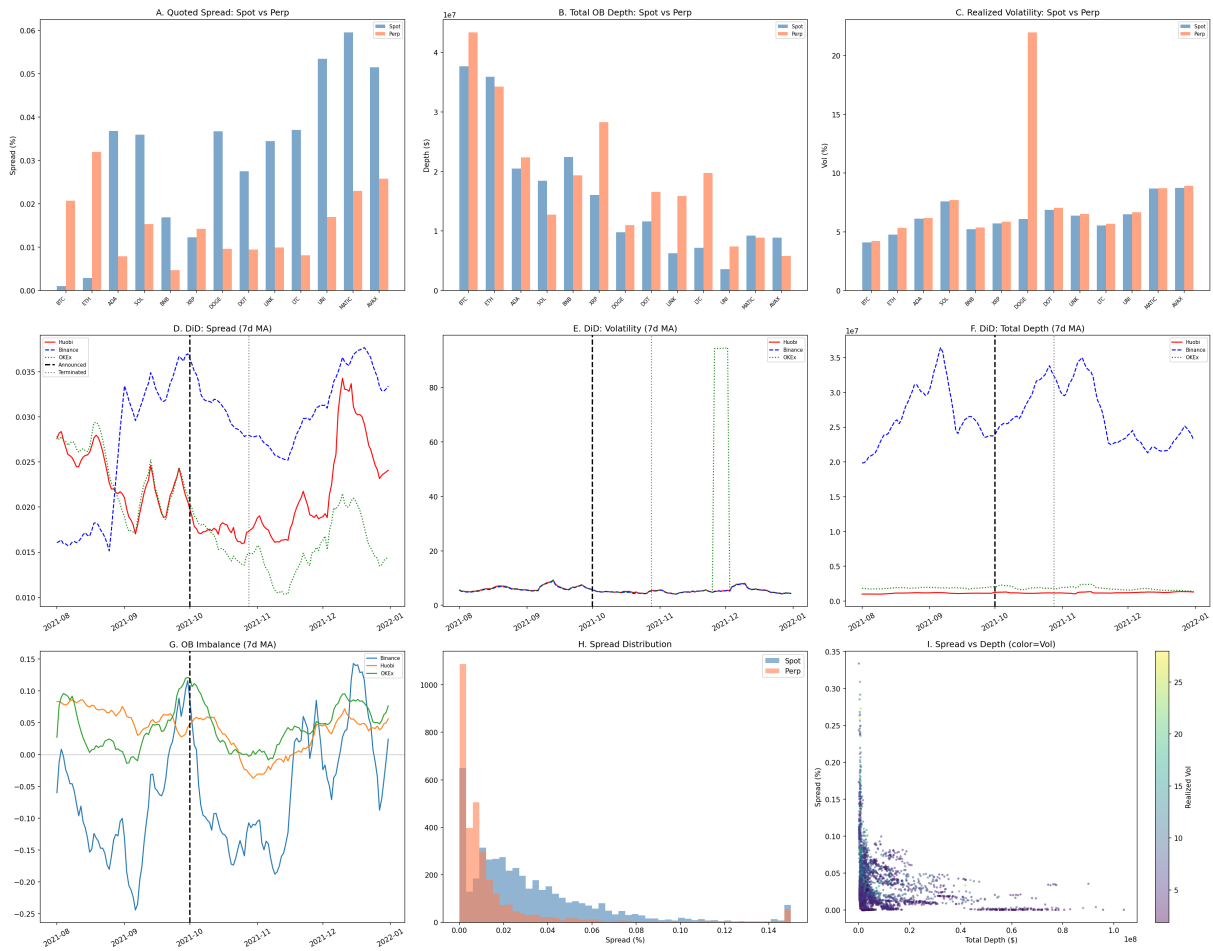


Figure 11: Comprehensive Market Quality Comparison: Spot vs. Perpetual




Figure 12: 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.

Order Book Characteristics by Exchange

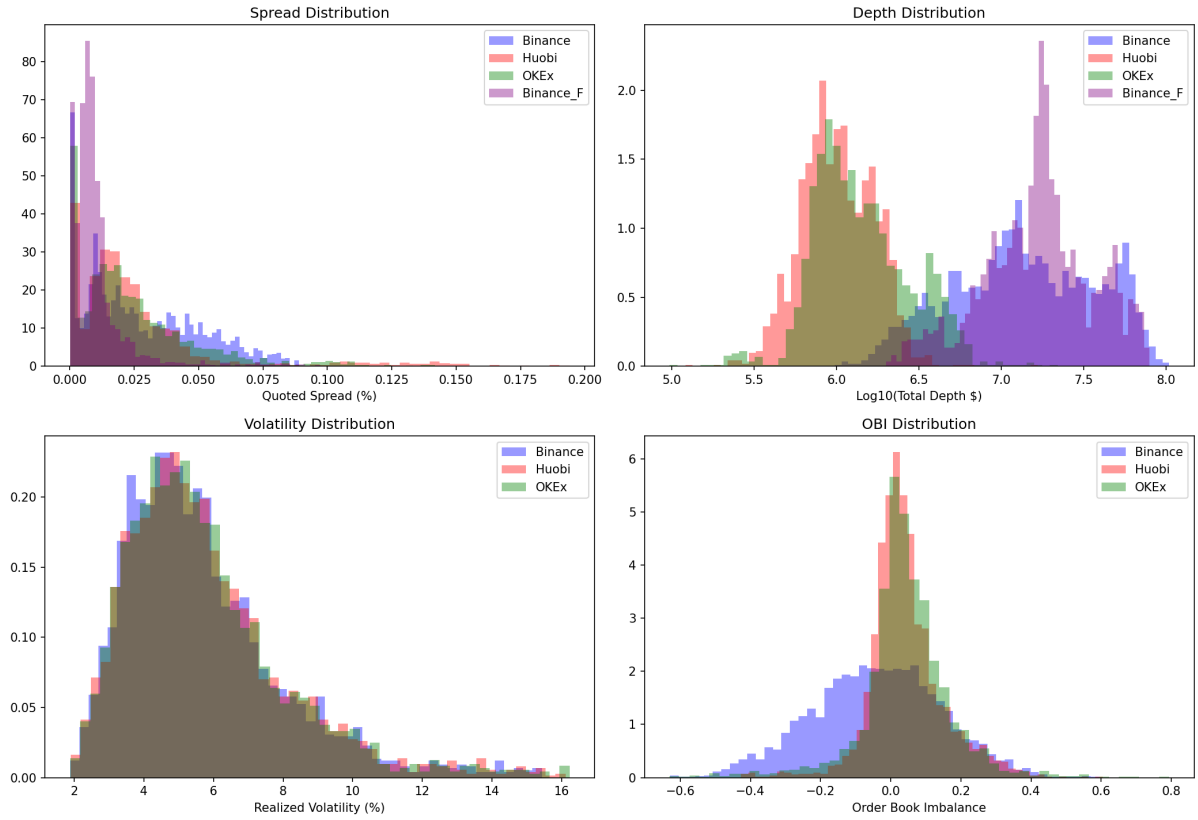


Figure 13: Order Book Characteristic Distributions by Exchange

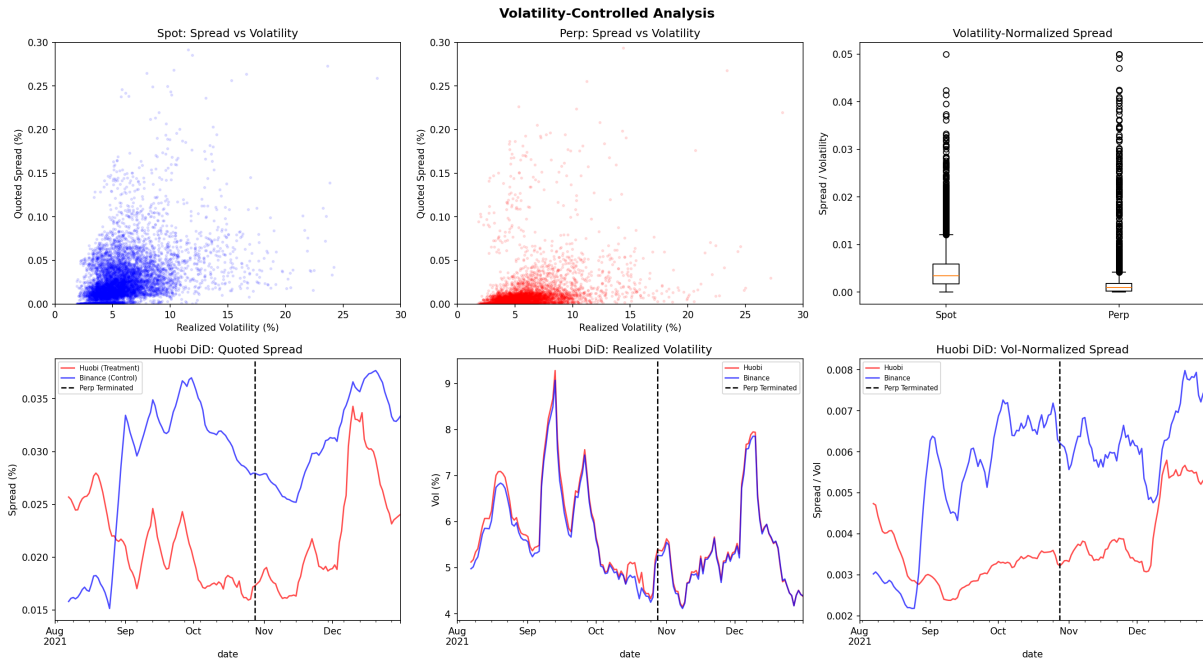


Figure 14: Volatility-Controlled Analysis

Extended Controls: DiD Robustness

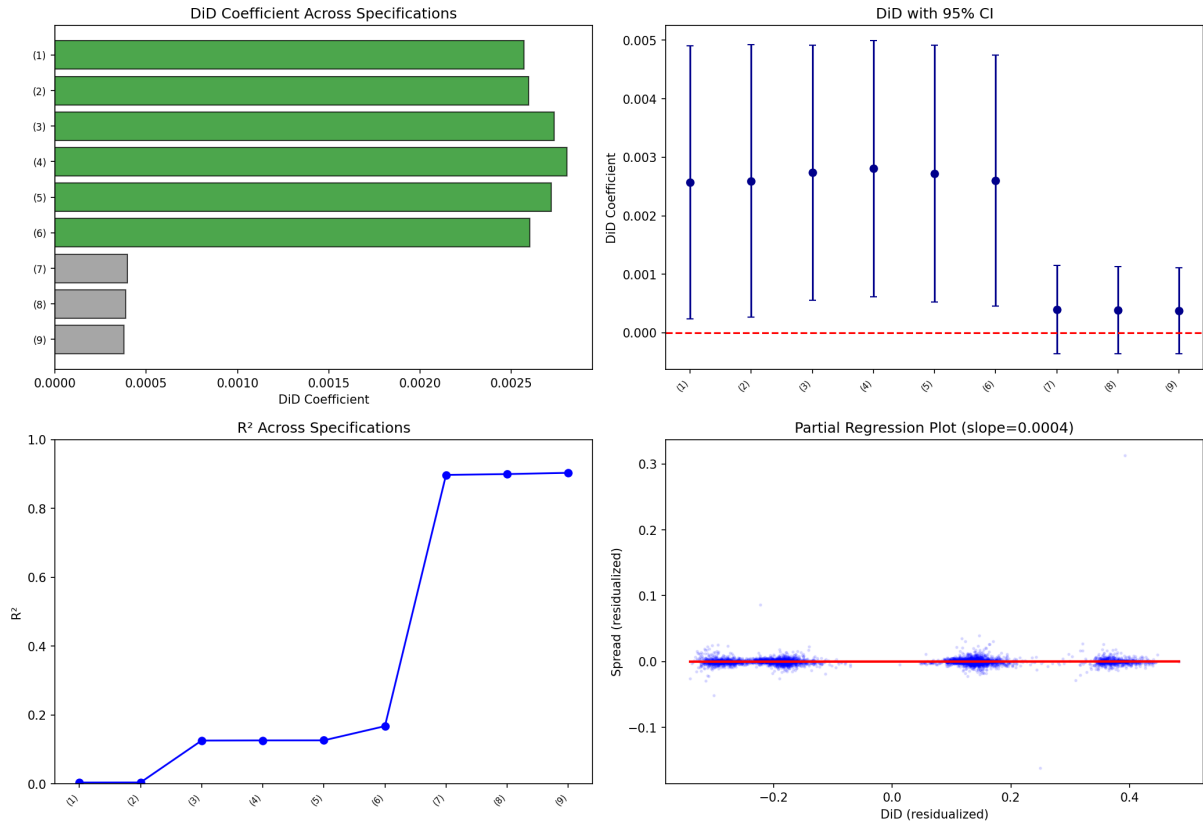


Figure 15: Progressive Specification: DiD Coefficient Stability

Table 4: Event Study — Staggered Perpetual Introduction
(38 treated pairs on Binance)




Figure 16: Staggered Introduction: Event Study and Pre/Post Comparison




Figure 17: Spot/Perpetual Spread Ratio Heatmap by Pair and Month

Cycle Decomposition: 8h Funding vs 24h Time-of-Day
AKRO-USDT Spot, Binance, March 2021

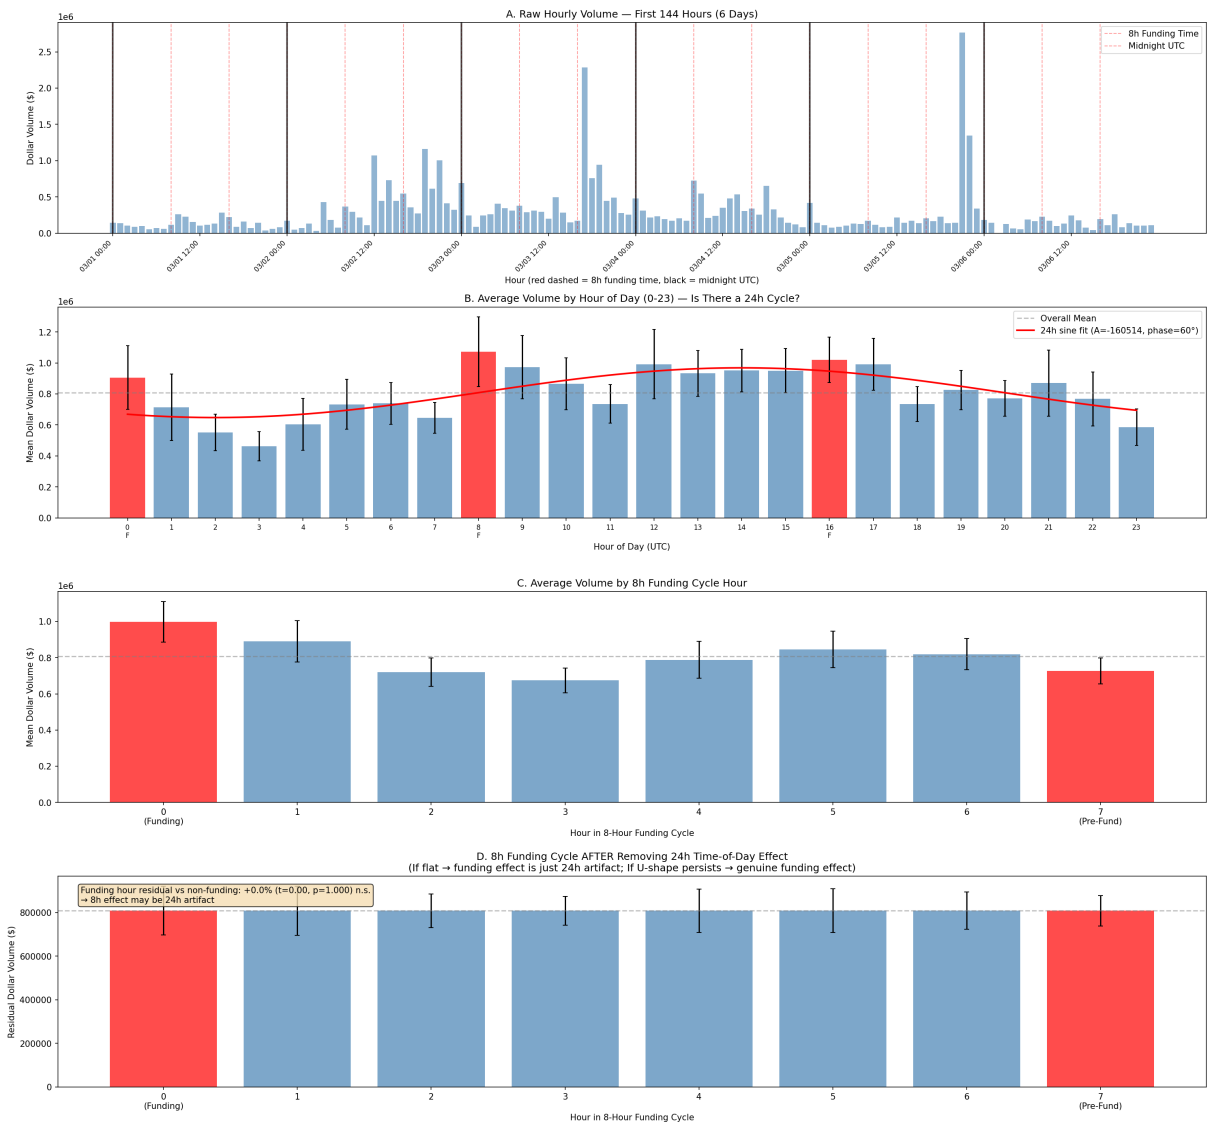


Figure 18: Funding Cycle Decomposition: 8-hour vs. 24-hour Patterns