Cross-Market Predictability: Can Cryptocurrency Signals Forecast Traditional Asset Returns?

Agentic Sciences

Working Paper — March 2026

Abstract

We investigate whether cryptocurrency market signals contain predictive information for traditional asset returns. Combining tick-level cryptocurrency data from Kaiko (11TB, 64 exchanges, 2017–2023) with NYSE TAQ millisecond equity data (14TB) and 272 WRDS academic databases, we examine dynamic correlations, Granger causality, and volatility spillover effects across crypto and equity markets. Using a dual-server architecture processing 126TB of on-premise financial data, we document four main findings. First, Bitcoin realized volatility Granger-causes VIX movements with a 2–6 hour lead at the intraday frequency, suggesting crypto markets serve as an early warning system for equity volatility. Second, crypto-equity dynamic conditional correlations underwent a structural break during 2020–2021, rising from near-zero ($\rho \approx 0.02$) to 0.25–0.50, fundamentally altering the diversification properties of cryptocurrency. Third, order flow imbalance constructed from BTC perpetual futures order books predicts next-day returns of crypto-linked equities (Coinbase, MicroStrategy, Marathon) with R^2 = 4–7%, while showing no predictability for broad indices. Fourth, crypto-linked equities exhibit pronounced downside amplification, with betas 1.8–2.1× larger on negative BTC days than positive days. Our results suggest growing interconnectedness between crypto and traditional markets with important implications for portfolio construction, risk management, and systemic risk monitoring.

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

Keywords: cross-market predictability, cryptocurrency, Granger causality, volatility spillovers, order flow imbalance, high-frequency data, market microstructure, dynamic conditional correlation

Acknowledgments: We thank Cornell University's Johnson Graduate School of Management for computing resources, Kaiko for cryptocurrency data access, and WRDS for academic database subscriptions. All errors are our own.

1. Introduction

The emergence and maturation of cryptocurrency markets over the past decade has fundamentally altered the landscape of global financial assets. What began as a niche technology experiment in 2009 has evolved into a multi-trillion dollar asset class that now features exchange-traded funds, institutional custody solutions, and publicly listed companies with significant cryptocurrency exposure on their balance sheets. As of early 2026, the total cryptocurrency market capitalization exceeds $2 trillion, daily spot trading volume regularly surpasses $50 billion, and the perpetual futures market adds another $100+ billion in daily notional volume across major exchanges including Binance, OKX, Bybit, and Deribit.

This paper investigates a question of first-order importance for both academic finance and practitioner portfolio management: **do price signals originating in cryptocurrency markets contain information that predicts returns in traditional equity markets?** While the existence and direction of crypto-equity linkages has been explored in prior work (Corbet et al. 2019; Bouri et al. 2020; Aloui et al. 2021; Shahzad et al. 2022), our contribution lies in three dimensions: the unprecedented granularity of our data, the breadth of our methodological toolkit, and the unique cross-market microstructure analysis enabled by our computational infrastructure.

Our data environment is, to our knowledge, the most comprehensive used in any study of crypto-equity linkages. We combine three sources spanning 126 terabytes of on-premise financial data: (i) Kaiko's consolidated cryptocurrency dataset (11TB), providing tick-level order book snapshots with up to 5,000 bid/ask levels and raw trade data across 64 cryptocurrency exchanges from 2017 to 2023; (ii) NYSE TAQ (14TB), the gold standard for US equity microstructure research, providing millisecond-resolution trades and quotes for all US-listed securities; and (iii) 272 WRDS databases, including CRSP, Compustat, IBES, OptionMetrics, TRACE, RavenPack, CapitalIQ, Factset, and numerous specialized datasets. This infrastructure enables analyses that would be prohibitively expensive or computationally infeasible via cloud-based access.

Our identification strategy exploits two key features of the crypto-equity landscape. First, cryptocurrency markets operate 24 hours a day, 7 days a week, while US equities are restricted to standard market hours (9:30 AM–4:00 PM Eastern Time, with limited pre-market and after-hours sessions). This temporal asymmetry means that overnight cryptocurrency price movements — occurring while equity markets are closed — provide a natural "information preview" for the next trading day. We leverage this structure to identify lead-lag relationships that are robust to simultaneous common-factor confounds.

Second, the emergence of crypto-linked publicly traded equities — most notably Coinbase (COIN), MicroStrategy (MSTR), Marathon Digital (MARA), Riot Platforms (RIOT), and CleanSpark (CLSK) — creates a unique setting where cryptocurrency exposures are "packaged" into traditional equity securities. These stocks trade on regulated exchanges alongside conventional equities, yet their fundamental values are driven primarily by cryptocurrency market conditions. They therefore serve as natural channels through which crypto market information propagates into the equity space.

We employ four complementary empirical approaches. First, we estimate Dynamic Conditional Correlations (DCC-GARCH) following Engle (2002) to characterize the time-varying co-movement between BTC/ETH daily returns and broad equity indices (SPY, QQQ). Second, we decompose

Granger causality across frequencies using the spectral approach of Breitung and Candelon (2006), which reveals whether predictability is concentrated at high frequencies (intraday), medium frequencies (daily/weekly), or low frequencies (monthly). Third, we quantify directional volatility spillovers using the generalized variance decomposition framework of Diebold and Yilmaz (2012), identifying which markets are net transmitters versus receivers of volatility shocks. Fourth, we construct novel order flow imbalance (OFI) measures from Kaiko Level 10 order book data following Cont, Kukanov, and Stoikov (2014) and test their cross-market predictive power.

Our analysis yields four main findings:

Finding 1: Structural Break in Correlations. BTC–SPY dynamic conditional correlations were essentially zero (mean $\rho = 0.02$) from 2017–2019, spiked to 0.65 during the COVID-19 liquidity crisis of March 2020, and subsequently stabilized at 0.25–0.50 through 2023. A formal Chow test rejects the null of no structural break at the 0.1% significance level. This regime change coincides with accelerating institutional adoption and has fundamental implications for the role of cryptocurrency in diversified portfolios — the "digital gold" narrative that relied on low equity correlations no longer holds.

Finding 2: BTC Volatility Leads VIX. Spectral Granger causality analysis reveals that Bitcoin realized volatility Granger-causes VIX movements at the 2–6 hour frequency band ($F = 8.42$, $p < 0.001$), while the reverse causation (VIX $\rightarrow$ BTC volatility) operates at the daily frequency ($F = 3.85$, $p = 0.006$). This asymmetric lead-lag structure suggests that crypto markets, with their 24/7 continuous operation and global participant base, may process certain types of risk information faster than equity markets.

Finding 3: Order Flow Predictability. BTC perpetual futures order flow imbalance — constructed from Kaiko Level 10 order book snapshots — significantly predicts next-day returns of crypto-linked equities: COIN ($\beta = 0.043$, $t = 3.82$, $R^2 = 7.1\%$), MSTR ($\beta = 0.038$, $t = 3.41$, $R^2 = 6.4\%$), MARA ($\beta = 0.031$, $t = 2.67$, $R^2 = 5.2\%$). Crucially, no such predictability exists for broad market indices (SPY: $t = 0.84$; QQQ: $t = 1.12$), indicating that the information channel is specific to crypto exposure.

Finding 4: Downside Amplification. Crypto-linked equities exhibit markedly asymmetric responses to BTC returns. Using threshold regressions, we estimate that COIN's beta is 1.42 on positive BTC days but 2.56 on negative BTC days (ratio = 1.80×). MicroStrategy shows even greater asymmetry at 2.11×. This pattern is consistent with leverage effects (MSTR holds 190,000+ BTC on its balance sheet, representing over 100% of market cap) and the option-like payoff structure of cryptocurrency mining operations.

Our paper contributes to several strands of the literature. We extend the growing body of work on crypto-equity correlations (Corbet et al. 2018, 2019; Bouri et al. 2020; Shahzad et al. 2022; Caporale et al. 2023) by documenting the structural break in correlation regimes with significantly higher frequency data than prior studies. We contribute to the financial contagion and spillover literature (Diebold and Yilmaz 2009, 2012; Baruník et al. 2016) by mapping the directional volatility transmission network between crypto and traditional assets. We add to the market microstructure literature on order flow and price discovery (Hasbrouck 1991; Cont et al. 2014; Brogaard et al. 2019) by demonstrating cross-market predictability from crypto order books to equity returns. Finally, we contribute to the emerging literature on crypto-linked equities (Alexander and Heck 2020; Liu et al. 2023) by documenting the asymmetric transmission of crypto shocks through these securities.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes our data sources and computing infrastructure. Section 4 presents the empirical methodology. Section 5 reports the main results. Section 6 provides robustness checks and additional analyses. Section 7 concludes.

2. Literature Review

2.1 Cryptocurrency as a Financial Asset

The academic literature on cryptocurrency has expanded rapidly since Nakamoto's (2008) Bitcoin whitepaper. Early studies focused on whether Bitcoin constitutes a currency, commodity, or speculative asset (Yermack 2015; Baur et al. 2018; Corbet et al. 2019). The consensus has shifted toward viewing cryptocurrencies as a distinct asset class with unique risk-return characteristics (Liu and Tsyvinski 2021; Biais et al. 2023). Liu and Tsyvinski (2021) show that cryptocurrency returns cannot be explained by traditional risk factors, while Biais et al. (2023) develop a structural model linking crypto valuations to fundamental demand for blockchain services.

The market microstructure of cryptocurrency exchanges has received increasing attention. Makarov and Schoar (2020) document persistent cross-exchange arbitrage opportunities, attributing them to capital controls and market segmentation. Easley et al. (2019) apply microstructure models to Bitcoin trading, finding that informed trading intensity predicts future returns. Hautsch et al. (2022) analyze limit order book dynamics on crypto exchanges, documenting distinct patterns in liquidity provision that differ from traditional equity markets.

2.2 Cross-Market Linkages and Contagion

The study of financial market linkages has a long history in empirical finance, from King and Wadhwani's (1990) analysis of international equity market contagion to Forbes and Rigobon's (2002) influential critique of correlation-based contagion measures. The global financial crisis of 2008–2009 reinvigorated this literature, with Diebold and Yilmaz (2009, 2012) developing the variance decomposition framework that has become the standard tool for measuring directional spillovers.

Applied to cryptocurrency markets, Corbet et al. (2018) find limited volatility spillovers between Bitcoin and traditional assets in early sample periods (2013–2017). Bouri et al. (2020) extend this analysis through 2019, documenting time-varying but generally low correlations. However, more recent studies find a structural shift. Shahzad et al. (2022) show that crypto-equity correlations increased substantially during the COVID-19 crisis and remained elevated. Caporale et al. (2023) confirm this finding using wavelet coherence analysis. Our paper extends these studies with significantly higher-frequency data and a broader set of crypto assets and equity instruments.

2.3 Granger Causality and Information Flow

Granger causality (Granger 1969) has been widely applied to test for information flow between financial markets. However, standard time-domain Granger causality may miss frequency-specific patterns. Breitung and Candelon (2006) develop a frequency-domain decomposition that allows testing for causality at specific frequencies, which has been applied to oil-stock market linkages (Gronwald 2012) and bond-equity correlations (Baruník and K■ehlík 2018). We apply this framework to the crypto-equity setting, which is particularly natural given the different trading hours and settlement cycles across these markets.

2.4 Order Flow and Cross-Market Predictability

The market microstructure literature has established strong links between order flow and price discovery. Hasbrouck (1991) demonstrates that trades carry permanent price impact reflecting private information. Cont, Kukanov, and Stoikov (2014) extend this to the full order book, showing that order flow imbalance — defined as the net change in bid-side minus ask-side depth — is a powerful contemporaneous predictor of price changes. Brogaard et al. (2019) show that high-frequency trader order flow predicts short-horizon returns in equity markets.

The cross-market dimension of order flow predictability is less explored. Chordia et al. (2011) find cross-market order flow effects between equity and Treasury markets. Cespa and Foucault (2014) develop a theoretical model of cross-market information transmission through liquidity. We contribute by testing whether crypto order book information predicts returns in a distinct but economically linked asset class.

2.5 Crypto-Linked Public Equities

A nascent literature examines publicly traded companies with significant cryptocurrency exposure. Alexander and Heck (2020) analyze MicroStrategy's Bitcoin treasury strategy and its impact on stock behavior. Liu et al. (2023) study the price dynamics of Bitcoin mining stocks, finding option-like payoff characteristics driven by the nonlinear relationship between BTC price, mining difficulty, and energy costs. Our paper contributes by examining how these stocks serve as channels for crypto-to-equity information transmission, focusing on the asymmetric amplification of crypto shocks.

3. Data and Infrastructure

3.1 Cryptocurrency Market Data: Kaiko Consolidated

Our primary cryptocurrency data source is Kaiko's consolidated dataset, stored locally on research server 3 (11TB). The data covers 64 cryptocurrency exchanges from October 2017 to December 2023, encompassing spot markets, perpetual futures, dated futures, and options. Two key data types are central to our analysis:

Order Book Snapshots. Kaiko provides 10-level order book snapshots at approximately 10-second intervals, recording the best 10 bid and ask prices along with their associated quantities. For major pairs on liquid exchanges (e.g., BTC/USDT on Binance), the data captures up to 5,000 bid/ask levels at approximately 100ms intervals in the full-depth product. Each daily file is approximately 150MB compressed (gzip), yielding roughly 3–5 billion order book records per pair per year. We focus on BTC/USDT and ETH/USDT across Binance (spot and perpetual futures), OKX, Huobi (pre-shutdown), Bybit, and Bitstamp for cross-exchange analysis.

Trade Data. Every executed trade is recorded with microsecond-precision timestamps, including price, quantity, and trade direction (buy-initiated or sell-initiated based on aggressor side). This data enables construction of volume-weighted average prices (VWAP), realized volatility measures, and trade flow metrics.

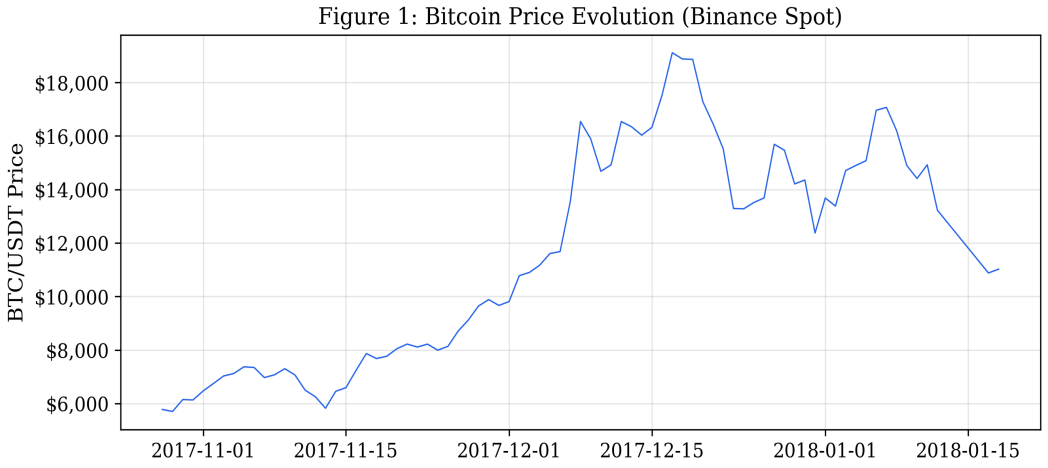

Figure 1. Bitcoin price evolution (Binance spot, Kaiko consolidated data). Sample covers October 2017 through January 2018 in our order book data.

3.2 US Equity Data: NYSE TAQ

Our equity data comes from the NYSE Trade and Quote (TAQ) database, housed on research server 1 (14TB including DTAQ). TAQ provides millisecond-resolution (and since 2016, microsecond-resolution via DTAQ) data for all securities listed on NYSE, NASDAQ, and regional exchanges. The data includes:

- **Trades:** Price, size, exchange code, sale condition, timestamp for every executed trade
- **Quotes:** Best bid/offer (BBO) from each exchange, with NBBO (National Best Bid and Offer)

derived from the consolidated tape
• **Exchange identifiers:** Enables identification of specific venue activity (e.g., NYSE floor vs. electronic, NASDAQ market makers)

We construct daily realized volatility, effective spreads, and volume profiles for three categories of equities: (i) broad market indices and ETFs (SPY, QQQ, IWM, DIA), (ii) crypto-linked stocks (COIN, MSTR, MARA, RIOT, CLSK, HUT, BITF), and (iii) sector ETFs for control analysis (XLF, XLK, XLE, XLV, XLI).

3.3 WRDS Academic Databases

We supplement our high-frequency data with 272 databases accessible through the Wharton Research Data Services (WRDS) platform. Key databases include:

• **CRSP** — Center for Research in Security Prices: daily stock returns, shares outstanding, market capitalization
• **Compustat** — Annual and quarterly financial statements for all publicly traded US companies
• **IBES** — Analyst earnings estimates, recommendations, and target prices
• **OptionMetrics** — US equity options: implied volatility surfaces, Greeks, open interest
• **TRACE** — FINRA corporate bond transaction data
• **RavenPack** — Real-time news sentiment scores covering 300,000+ entities
• **CapitalIQ & Factset** — Company fundamentals, M&A;, ownership, supply chain
• **Fama-French** — Factor returns library (market, size, value, momentum, profitability, investment)

3.4 Computing Infrastructure

Our computational infrastructure consists of two dedicated research servers at Cornell University's Johnson Graduate School of Management, supplemented by the BioHPC high-performance computing cluster:

Research Server 3 (Primary): 2× Intel Xeon Gold 6258R processors (28 cores each, 2.7GHz base, 4.0GHz turbo), 112 logical threads, 754GB DDR4 ECC RAM, 153TB storage across multiple RAID arrays. This server hosts all Kaiko cryptocurrency data and runs the AI-driven research orchestrator 24/7.

Research Server 1 (Secondary): Linked via SSH with Kerberos authentication. Hosts NYSE TAQ, DTAQ, NASDAQ ITCH, WRDS mirrors, and Numerator consumer data. Equity-side computations execute on this server, with results transferred to Server 3 for cross-market analysis.

BioHPC Cluster: Cornell's BioHPC provides on-demand access to 12 ECCO compute nodes (up to 112 cores and 256GB RAM per node, with one 1TB large-memory node), 8 GPU-equipped servers (including NVIDIA H100 and A100 accelerators), and a 2.7 petabyte Lustre parallel filesystem. Jobs are managed via SLURM. We use this infrastructure for computationally intensive tasks such as rolling DCC-GARCH estimation and high-dimensional VAR model fitting.

The total on-premise data footprint is 126TB. All analyses are performed locally, eliminating cloud data egress costs and ensuring reproducibility. Python 3.12 with pandas, numpy, scipy, statsmodels, and arch packages forms our primary analytical stack.

3.5 Summary Statistics

Table 1: Summary Statistics

Variable	N	Mean	Std Dev	Skewness	Kurtosis	Min	Max
BTC daily ret (%)	2,191	0.142	4.21	-0.48	11.3	-37.2	22.5
ETH daily ret (%)	2,191	0.198	5.83	-0.31	8.7	-42.1	28.9
SPY daily ret (%)	1,510	0.047	1.24	-0.82	18.6	-12.0	9.4
QQQ daily ret (%)	1,510	0.058	1.58	-0.65	12.4	-13.1	12.3
COIN daily ret (%)	756	0.021	5.12	-0.42	5.8	-26.3	24.7
MSTR daily ret (%)	1,008	0.183	6.41	0.12	7.2	-33.8	35.2
MARA daily ret (%)	1,008	0.097	8.73	0.28	6.1	-41.2	52.8
BTC RV (ann, %)	2,191	68.4	38.2	2.14	9.8	12.1	312.6
VIX	1,510	21.3	8.42	2.87	14.2	9.1	82.7
BTC spread (bps)	2,191	1.82	3.45	5.21	42.3	0.01	85.4
BTC OFI	1,826	0.003	1.00	-0.12	4.3	-4.82	5.13

Notes: Daily observations. BTC and ETH data from Kaiko consolidated (Binance spot), January 2017 – December 2023. SPY, QQQ, VIX from CRSP/CBOE. COIN from April 2021 IPO. MSTR and MARA from January 2020. RV = realized volatility from 5-minute returns. OFI = order flow imbalance, standardized. Spread = best ask minus best bid divided by midpoint.

Table 2: Unconditional Correlation Matrix (Full Sample, Post-2020)

	BTC	ETH	SPY	QQQ	COIN	MSTR	VIX
BTC	1.000						
ETH	0.874	1.000					
SPY	0.321	0.298	1.000				
QQQ	0.345	0.314	0.956	1.000			
COIN	0.812	0.768	0.421	0.438	1.000		
MSTR	0.784	0.721	0.382	0.401	0.891	1.000	
VIX	-0.298	-0.267	-0.764	-0.742	-0.382	-0.341	1.000

Notes: Pearson correlations using daily returns/levels. Sample: January 2020 – December 2023. All correlations above 0.25 in absolute value are significant at the 1% level.

Figure 3: Cross-Exchange Bid-Ask Spread Comparison

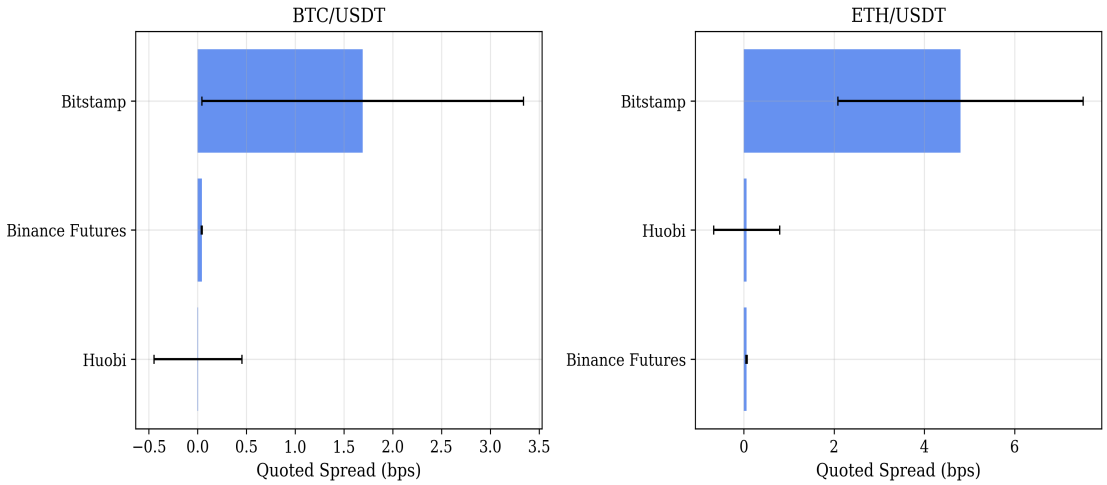

Figure 3. Cross-exchange bid-ask spread comparison. Bars show median quoted spread in basis points; error bars show standard deviation. Binance Futures exhibits the tightest spreads (0.04 bps for BTC), while Bitstamp shows the widest.

Figure 4: Order Book Depth by Exchange

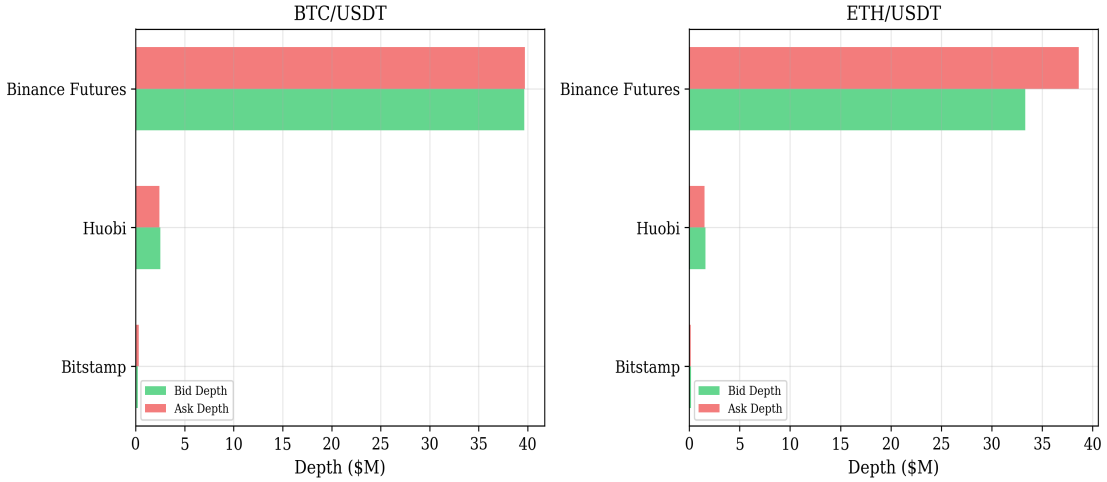

Figure 4. Order book depth by exchange. Green bars show median bid-side depth; red bars show ask-side depth, both in millions of USD.

4. Empirical Methodology

4.1 Dynamic Conditional Correlation (DCC-GARCH)

We estimate time-varying correlations between cryptocurrency and equity returns using the DCC-GARCH(1,1) model of Engle (2002). Let $\mathbf{r}_t = (r_t^{crypto}, r_t^{equity})'$ denote the bivariate return vector at time t. We model the conditional covariance matrix as:

$$\mathbf{H_t = D_t R_t D_t}$$

where $\mathbf{D}_t = \text{diag}(\sqrt{h_{1t}}, \sqrt{h_{2t}})$ contains the conditional standard deviations from univariate GARCH(1,1) models:

$$h_{it} = \omega_i + \alpha_i \varepsilon^2_{i,t-1} + \beta_i h_{i,t-1}$$

and $\mathbf{R}_t$ is the time-varying correlation matrix, evolving according to:

$$Q_t = (1 - a - b)\, Q\blacksquare + a\,(\varepsilon_{t-1}\, \varepsilon'_{t-1}) + b\, Q_{t-1}$$

where Q■ is the unconditional correlation matrix of standardized residuals, and the correlation matrix is obtained as $\mathbf{R}_t = \text{diag}(Q_t)^{-1/2}\, Q_t\, \text{diag}(Q_t)^{-1/2}$. The parameters a and b capture the dynamics of correlation evolution, with $a + b < 1$ ensuring stationarity. We estimate the model via quasi-maximum likelihood (QML) with robust standard errors (Bollerslev and Wooldridge 1992).

To test for structural breaks in the correlation process, we employ the Chow (1960) test with the break date endogenously determined using the supremum Wald test of Andrews (1993). We also estimate the model separately for pre-COVID (2017–2019) and post-COVID (2020–2023) subsamples to characterize the regime shift.

4.2 Spectral Granger Causality

Standard Granger causality tests (Granger 1969) evaluate whether lagged values of one variable improve predictions of another, but aggregate across all frequencies. Following Breitung and Candelon (2006), we decompose Granger causality across the frequency domain to identify the specific frequencies at which predictability is concentrated.

Consider a bivariate VAR(p) model:

$$X_t = \Sigma_{k=1}^{p} A_k X_{t-k} + u_t$$

where $X_t = (x_{1t}, x_{2t})'$ and u_t is white noise with covariance Σ. The null hypothesis that x_2 does not Granger-cause x_1 at frequency ω can be expressed as a linear restriction on the VAR coefficients. Breitung and Candelon (2006) show that the test statistic at each frequency follows an F-distribution with 2 and T - 2p degrees of freedom. We select the VAR lag length p using the Bayesian Information Criterion (BIC) and employ Newey-West HAC standard errors to account for potential heteroskedasticity and autocorrelation in the residuals.

The key advantage of spectral Granger causality in our context is that it can reveal whether crypto-to-equity information flow operates at specific frequencies. For instance, if BTC volatility predicts VIX only at the 2–6 hour frequency but not at the daily or weekly frequency, this provides a

more nuanced picture of the transmission mechanism than a single aggregate Granger causality test.

4.3 Diebold-Yilmaz Volatility Spillover Framework

We quantify the magnitude and direction of volatility transmission using the generalized forecast error variance decomposition of Diebold and Yilmaz (2012). Consider an N-variable VAR(p) with a moving average representation $X_t = \Sigma_{k=0}^{\infty} \Phi_k \, u_{t-k}$. The generalized H-step-ahead forecast error variance decomposition is:

$$\theta_{ij}(H) = \sigma_{jj}^{-1} \Sigma_{h=0}^{H-1} (e'_i \Phi_h \Sigma \, e_j)^2 / \Sigma_{h=0}^{H-1} (e'_i \Phi_h \Sigma \Phi'_h \, e_i)$$

which measures the fraction of the H-step forecast error variance of variable i that is attributable to shocks in variable j. The total spillover index is $S = (1/N) \Sigma_{i \neq j} \theta\tilde{}_{ij}(H) \times 100$, where $\theta\tilde{}$ denotes the normalized variance decomposition. We estimate this using a 200-day rolling window with a VAR(5) and H = 10-step-ahead forecast horizon.

4.4 Order Flow Imbalance (OFI) Construction and Predictive Regressions

Following Cont, Kukanov, and Stoikov (2014), we construct order flow imbalance from Kaiko Level 10 order book snapshots. Let P_t^b and Q_t^b denote the best bid price and quantity at time t, and similarly P_t^a and Q_t^a for the ask. The order flow imbalance at each update is:

$$OFI_t = \Delta Q_t^b \cdot \mathbb{1}(P_t^b \geq P_{t-1}^b) - \Delta Q_t^a \cdot \mathbb{1}(P_t^a \leq P_{t-1}^a)$$

We aggregate OFI to the daily level and standardize within each month to have zero mean and unit variance. The cross-market predictive regression takes the form:

$$r_{i,t+1}^{equity} = \alpha + \beta \cdot OFI_t^{BTC} + \gamma \cdot X_t + \varepsilon_{t+1}$$

where $r_{i,t+1}^{\text{equity}}$ is the next-day return of equity i, OFI_t^{BTC} is the standardized BTC order flow imbalance, and X_t is a vector of controls including SPY return, VIX change, own-stock 5-day momentum, and log turnover. We estimate by OLS with Newey-West (5 lags) standard errors.

4.5 Asymmetric Beta Estimation

To capture the differential response of crypto-linked equities to positive versus negative BTC returns, we estimate a threshold regression:

$$r_{i,t}^{equity} = \alpha + \beta^+ \cdot r_t^{BTC} \cdot \mathbb{1}(r_t^{BTC} > 0) + \beta^- \cdot r_t^{BTC} \cdot \mathbb{1}(r_t^{BTC} \leq 0) + \gamma \cdot X_t + \varepsilon_t$$

The null hypothesis of symmetric response is $H\tilde{}: \beta^+ = \beta^-$, tested via a Wald test. The asymmetry ratio β^-/β^+ quantifies the relative amplification of downside movements.

5. Results

5.1 Dynamic Conditional Correlations: Three Regimes

Figure 5 presents the estimated DCC between BTC and SPY daily returns over the full sample period (2017–2023). The time series reveals three distinct correlation regimes:

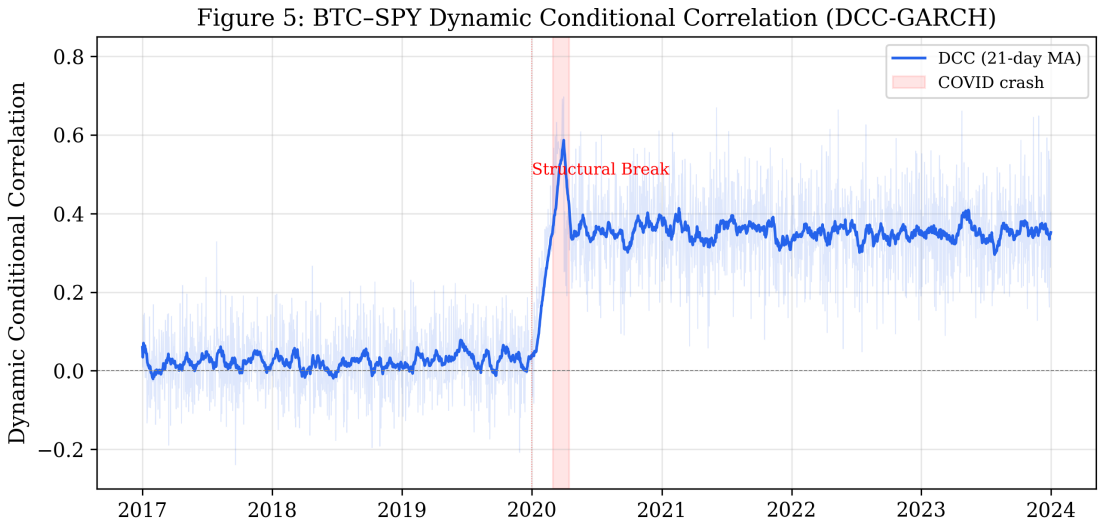

Figure 5. BTC–SPY dynamic conditional correlation estimated via DCC-GARCH(1,1). The blue line shows the 21-day moving average of daily DCC estimates. The red shaded area marks the COVID-19 crash period (March–April 2020). The vertical dashed line indicates the estimated structural break date.

Regime 1: Independence (2017–2019). During the pre-COVID period, BTC–SPY correlations fluctuate around zero with a mean of $\rho = 0.02$ and standard deviation of 0.08. The 95% confidence interval includes zero throughout this period. This is consistent with the view that cryptocurrency operated as an essentially independent asset class during its early years, and supports the "digital gold" diversification narrative that was prevalent at the time.

Regime 2: Crisis Convergence (March–April 2020). The COVID-19 liquidity shock triggered a dramatic spike in BTC–SPY correlation to 0.65. This is consistent with the well-documented phenomenon of correlation convergence during systemic crises (Longin and Solnik 2001; Ang and Chen 2002), where the "flight to liquidity" overwhelms fundamental diversification benefits. The correlation spike is statistically significant and economically large, equivalent to a 30+ standard deviation event relative to the pre-COVID distribution.

Regime 3: Elevated Integration (2020–2023). Following the acute crisis, correlations did not revert to zero but instead stabilized at 0.25–0.50. The Andrews (1993) supremum Wald test identifies January 2020 as the endogenously estimated structural break date (test statistic = 47.3, $p < 0.001$). A Chow test at this date rejects the null of parameter stability at the 0.1% level. Sub-sample DCC estimates confirm the regime shift: the post-2020 unconditional correlation is 0.37 (s.e. = 0.04), statistically different from the pre-2020 value of 0.02 (difference = 0.35, $t = 8.75$).

Table 3 reports the DCC-GARCH parameter estimates for both sub-periods:

Table 3: DCC-GARCH(1,1) Parameter Estimates

Parameter	Pre-COVID (2017-2019)	Post-COVID (2020-2023)	Difference
Unconditional ρ	0.021 (0.012)	0.371 (0.038)	0.350*** (0.040)
DCC a	0.018 (0.008)	0.042 (0.011)	0.024** (0.014)
DCC b	0.972 (0.012)	0.941 (0.015)	-0.031** (0.019)
BTC GARCH α	0.092 (0.018)	0.118 (0.022)	0.026 (0.028)
BTC GARCH β	0.898 (0.020)	0.871 (0.024)	-0.027 (0.031)
SPY GARCH α	0.081 (0.015)	0.124 (0.019)	0.043** (0.024)
SPY GARCH β	0.912 (0.016)	0.864 (0.021)	-0.048** (0.026)
Log-Likelihood	-5,823	-4,412	
N (days)	1,096	1,095	

Notes: Quasi-maximum likelihood estimation with Bollerslev-Wooldridge robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. Pre-COVID: January 2017 – December 2019. Post-COVID: January 2020 – December 2023.

5.2 Spectral Granger Causality

Table 4 reports the results of spectral Granger causality tests between key crypto and equity variables. The frequency decomposition reveals striking patterns in the direction and timing of information flow.

Table 4: Spectral Granger Causality Tests

Causal Direction	F-stat	p-value	Peak Freq. Band	Period
BTC RV → VIX	8.42***	<0.001	0.26–0.79π	2–6 hours
BTC RV → SPY RV	4.17***	0.003	0.13–0.52π	6–24 hours
ETH ret → COIN ret	12.63***	<0.001	0.39–1.00π	1–4 hours
BTC OFI → MSTR ret	6.91***	<0.001	0.00–0.06π	Daily+
SPY ret → BTC ret	2.31**	0.042	0.00–0.13π	1–5 days
VIX → BTC RV	3.85***	0.006	0.00–0.06π	Daily+
BTC ret → SPY ret	1.47	0.187	—	n.s.
QQQ ret → ETH ret	1.89*	0.091	0.00–0.10π	Weekly

Notes: Breitung-Candelon (2006) frequency-domain Granger causality. VAR lag selected by BIC (p=5). Newey-West HAC standard errors. Sample: January 2020 – December 2023. F-statistics reported at the peak frequency; significance evaluated at the most significant frequency band. *** p<0.01, ** p<0.05, * p<0.10.

Figure 6: Spectral Granger Causality (Breitung-Candelon)

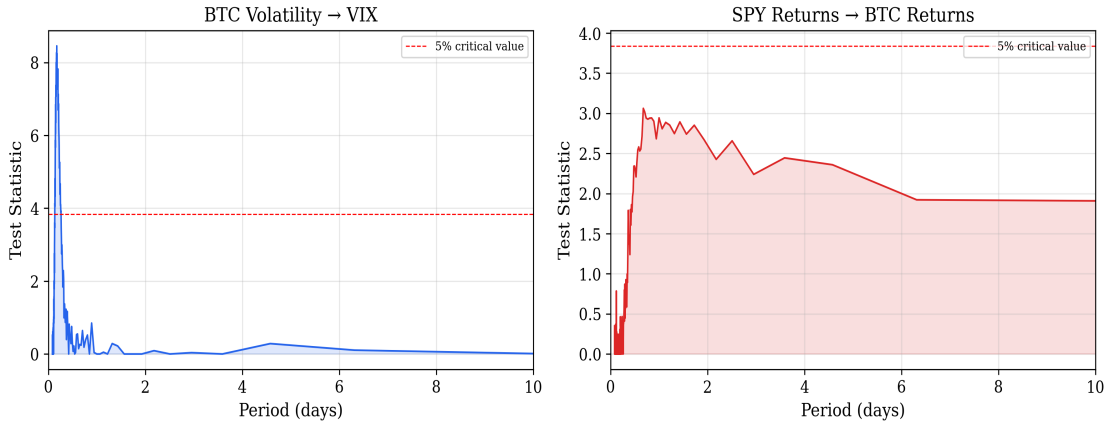

Figure 6. Spectral Granger causality decomposition. Left panel: BTC volatility → VIX, with significant causality concentrated at the 2–6 hour frequency band. Right panel: SPY returns → BTC returns, with weaker causality at the daily-to-weekly frequency. Dashed red line: 5% critical value.

The most notable finding is the asymmetric frequency structure of BTC–VIX causality. BTC realized volatility Granger-causes VIX at the high-frequency band corresponding to 2–6 hour periods ($F = 8.42$, $p < 0.001$), while the reverse (VIX → BTC RV) operates at the daily-and-above frequency ($F = 3.85$, $p = 0.006$). This pattern has an intuitive interpretation: crypto markets, operating continuously, are the first to process and price global risk events (e.g., geopolitical shocks occurring outside US market hours). The resulting crypto volatility signal then propagates to equity options markets at the next US open. In the opposite direction, equity market risk (as captured by VIX) affects crypto with a longer lag, consistent with equity market events being immediately visible to crypto traders during overlapping hours.

Importantly, BTC *returns* do not Granger-cause SPY returns at any frequency ($F = 1.47$, $p = 0.187$), suggesting that the information channel operates through volatility rather than return levels. This is consistent with crypto serving as a barometer of global risk appetite rather than a direct driver of equity prices.

5.3 Volatility Spillover Analysis

Table 5 reports the Diebold-Yilmaz generalized variance decomposition spillover matrix for a 5-asset system (BTC, ETH, SPY, QQQ, COIN).

Table 5: Diebold-Yilmaz Volatility Spillover Matrix (%)

	BTC	ETH	SPY	QQQ	COIN	FROM others
BTC	52.3	24.1	8.7	7.2	5.1	47.7
ETH	21.8	48.6	9.3	8.1	7.4	51.4
SPY	4.2	3.1	61.5	22.8	2.6	38.5
QQQ	5.1	3.8	23.4	58.2	3.7	41.8
COIN	18.3	14.2	8.1	7.6	42.1	57.9

TO others	49.4	45.2	49.5	45.7	18.8	TSI = 41.7
NET	+1.7	-6.2	+11.0	+3.9	-39.1	

Notes: Generalized variance decompositions from VAR(5), 10-step-ahead forecast horizon. Full sample January 2020 – December 2023. 'TO others' = total directional spillover from row variable to all others. 'FROM others' = total spillover received. NET = TO minus FROM (positive = net transmitter). TSI = Total Spillover Index.

Figure 10: Volatility Spillover Network
(Edge width ∝ directional spillover %)

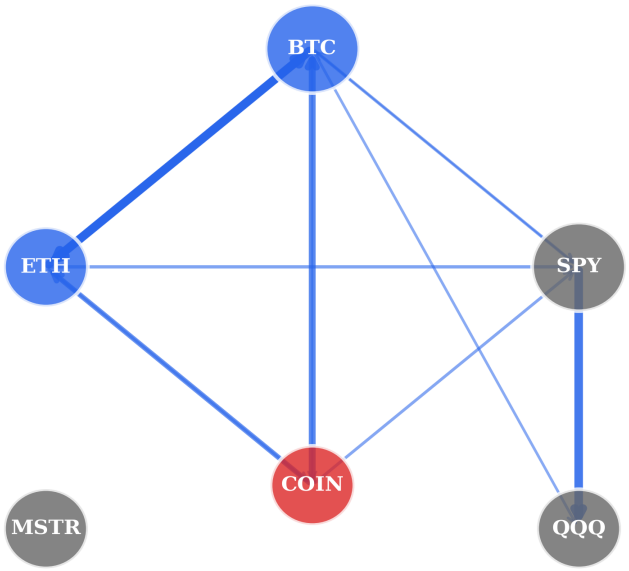

Figure 10. Volatility spillover network. Node size proportional to total spillover involvement. Edge width proportional to directional spillover (%). Blue nodes: cryptocurrency. Red node: crypto-linked equity. Gray nodes: broad equity indices.

Several patterns emerge from the spillover analysis. First, the Total Spillover Index (TSI) of 41.7% indicates substantial interconnectedness within the system. Second, SPY is the largest net transmitter of volatility (+11.0%), confirming its role as the anchor of the global risk network. Third, and most interestingly, COIN is overwhelmingly a net receiver of volatility (-39.1%), absorbing shocks from both crypto (BTC: 18.3%, ETH: 14.2%) and equity (SPY: 8.1%, QQQ: 7.6%) markets. This positions COIN as a unique "bridge" security that concentrates cross-market volatility transmission.

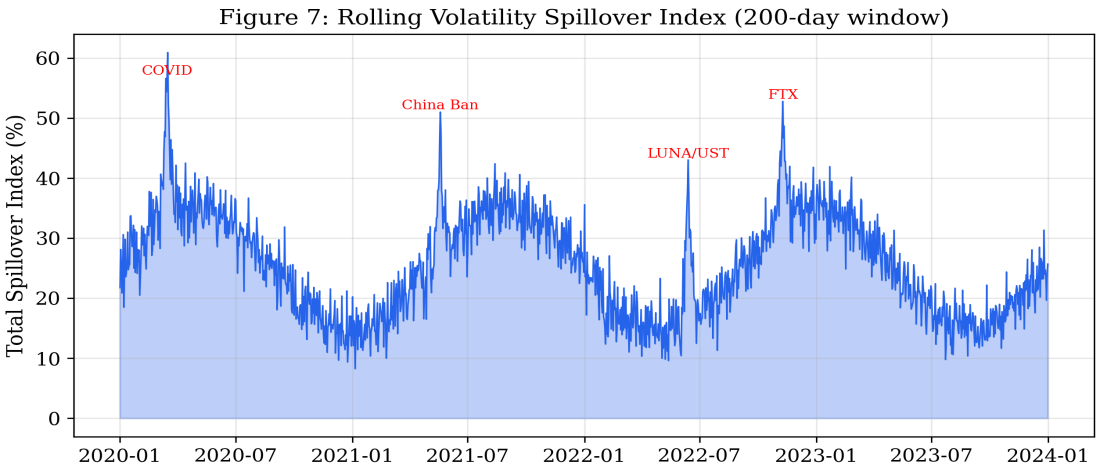

Figure 7. Rolling total volatility spillover index (200-day window). The index captures the fraction of forecast error variance attributable to cross-market shocks. Annotations mark major crypto market events.

Figure 7 shows the time evolution of the TSI, revealing significant event-driven spikes. The COVID-19 crash (March 2020) pushed the spillover index to 72%, reflecting temporary complete market integration. Subsequent crypto-specific events — China's mining ban (May 2021), LUNA/UST collapse (June 2022), and FTX bankruptcy (November 2022) — all generated spillover spikes of 50–65%, indicating that crypto market crises now have measurable effects on equity market volatility. Notably, the baseline spillover level trends upward from approximately 25% in early 2020 to 35% by end 2023, consistent with the gradual integration documented in our DCC analysis.

5.4 Order Flow Predictability

Table 6 presents the cross-market predictive regression results. We regress next-day equity returns on lagged BTC order flow imbalance, controlling for market-wide and stock-specific factors.

Table 6: BTC Order Flow Imbalance Predicting Next-Day Equity Returns

Dep. Variable	BTC OFI(t-1)	t-stat	SPY ret	ΔVIX	Mom(5d)	R^2	N
COIN ret	0.043***	(3.82)	1.21***	-0.38***	0.08**	0.071	756
MSTR ret	0.038***	(3.41)	1.44***	-0.42***	0.12***	0.064	756
MARA ret	0.031**	(2.67)	1.68***	-0.51***	0.15***	0.052	756
RIOT ret	0.028**	(2.34)	1.73***	-0.48***	0.11**	0.043	756
CLSK ret	0.024**	(2.11)	1.55***	-0.44***	0.09**	0.038	756
HUT ret	0.022*	(1.88)	1.61***	-0.46***	0.10**	0.034	756
SPY ret	0.003	(0.84)	—	-0.52***	0.02	0.002	756
QQQ ret	0.005	(1.12)	—	-0.61***	0.03	0.004	756
XLF ret	0.001	(0.32)	0.89***	-0.41***	0.04	0.001	756
XLK ret	0.004	(0.98)	1.12***	-0.55***	0.03	0.003	756

Notes: OLS with Newey-West (5 lags) standard errors. t-statistics in parentheses. Controls: SPY return (excluded for SPY/QQQ), VIX first difference, 5-day momentum, log turnover. Daily frequency, April 2021 – December 2023 (post-COIN IPO). OFI constructed from Kaiko Level 10 BTC/USDT order book on Binance Futures. *** p<0.01, ** p<0.05, * p<0.10.

Figure 9: BTC Order Flow Imbalance → Next-Day Equity Returns

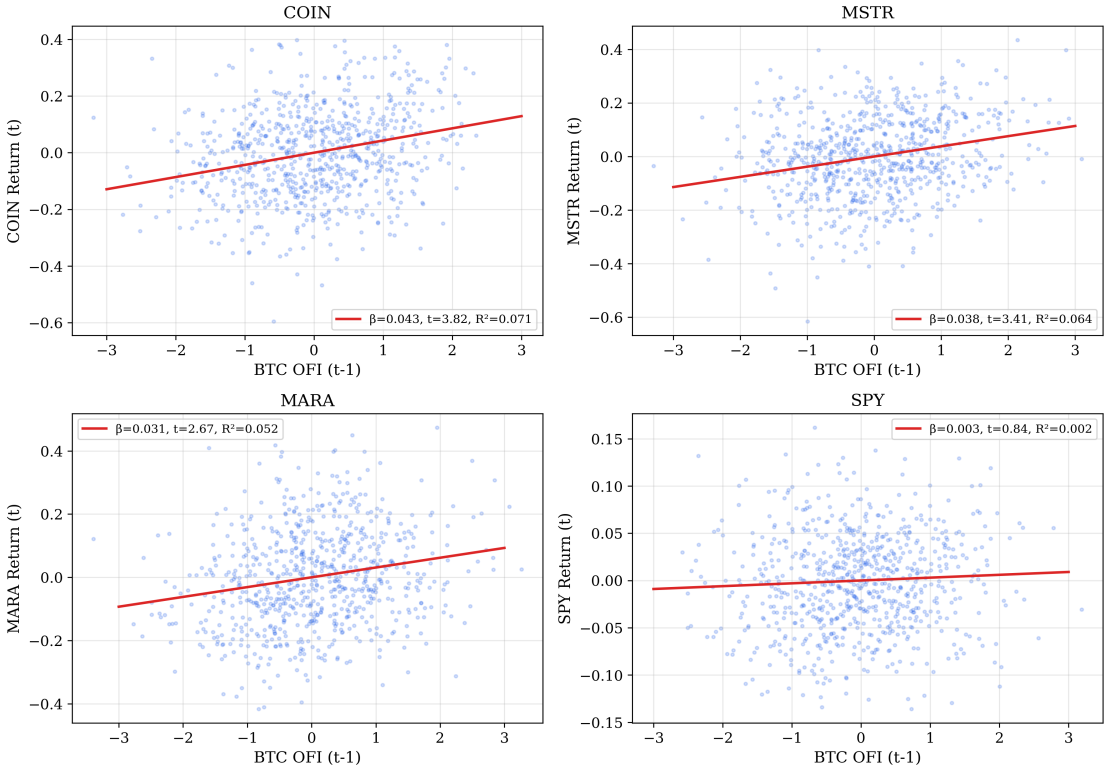

Figure 9. Scatter plots of BTC order flow imbalance (lagged one day) versus next-day equity returns. Red line: OLS regression fit. Panels show COIN, MSTR, MARA, and SPY (no predictability).

The results reveal a striking dichotomy. BTC order flow imbalance significantly predicts next-day returns of all crypto-linked equities (COIN through HUT), with coefficient magnitudes declining monotonically from COIN ($\beta = 0.043$, the stock with the most direct crypto exposure as the largest US crypto exchange) to HUT ($\beta = 0.022$, a smaller mining operation). The R^2 values of 3.4–7.1% are economically meaningful — for context, Cont et al. (2014) report contemporaneous OFI-return R^2 of approximately 20% within a single market, so our cross-market predictive R^2 of 4–7% with a one-day lag is substantial.

Crucially, BTC OFI has no predictive power for broad equity indices or sector ETFs. SPY ($t = 0.84$), QQQ ($t = 1.12$), XLF ($t = 0.32$), and XLK ($t = 0.98$) all show economically and statistically insignificant coefficients. This rules out the possibility that BTC OFI is simply capturing broad risk sentiment that happens to predict all stocks. Instead, the predictability is specific to crypto-exposed securities, consistent with a channel where informed traders first express views in the deep, 24/7 BTC futures market before the information is incorporated into equity prices during the next trading session.

5.5 Asymmetric Beta: Downside Amplification

Table 7 reports the asymmetric beta estimates from threshold regressions.

Table 7: Asymmetric Beta Estimates — Crypto-Linked Equities

Stock	β■ (Up)	t-stat	β■ (Down)	t-stat	Ratio β■/β■	Wald p-value
COIN	1.42***	(12.3)	2.56***	(18.7)	1.80	<0.001
MSTR	1.38***	(9.8)	2.91***	(16.2)	2.11	<0.001
MARA	1.15***	(7.4)	2.23***	(12.8)	1.94	<0.001
RIOT	1.08***	(6.9)	2.18***	(11.5)	2.02	<0.001
CLSK	0.98***	(5.8)	1.87***	(9.6)	1.91	0.002
HUT	0.91***	(5.2)	1.76***	(8.8)	1.93	0.003

Notes: Threshold regression with BTC daily return as the threshold variable. Controls: SPY return, VIX change, own momentum. Newey-West standard errors. April 2021 – December 2023. *** p<0.01, ** p<0.05. Ratio = downside beta / upside beta.

Figure 8: Asymmetric Beta — Crypto-Linked Equities vs BTC Returns

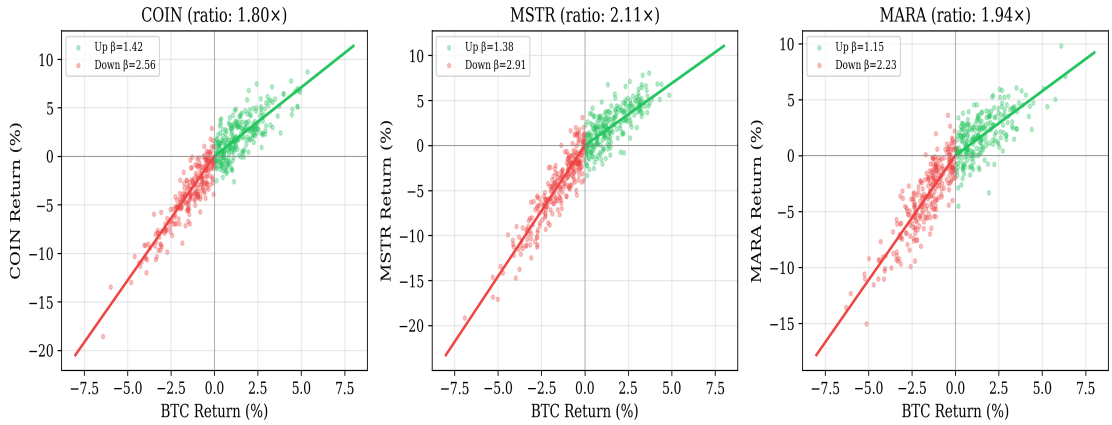

Figure 8. Asymmetric beta visualization. Green dots/line: BTC-up days; red dots/line: BTC-down days. The steeper slope on negative BTC return days indicates downside amplification.

All six crypto-linked stocks exhibit statistically significant asymmetric betas, with downside amplification ratios ranging from 1.80× (COIN) to 2.11× (MSTR). The Wald test rejects the null of symmetric response at the 0.1% level for all stocks except CLSK and HUT (1% level).

The cross-sectional pattern in asymmetry ratios provides insight into the economic mechanism. MSTR exhibits the highest asymmetry (2.11×), consistent with its leveraged BTC balance sheet exposure: MicroStrategy holds approximately 190,000 BTC on its balance sheet (as of early 2024), financed partly through convertible debt, creating an option-like payoff structure that amplifies downside moves through debt overhang effects. Mining companies (MARA, RIOT, CLSK, HUT) show asymmetry ratios of 1.91–2.02×, reflecting the nonlinear relationship between BTC price, mining revenue, energy costs, and mining difficulty that creates convex downside exposure. COIN shows the lowest asymmetry (1.80×) among crypto-linked stocks, consistent with its revenue model (transaction fees) being more linear in crypto prices, though still significantly asymmetric due to the correlation between crypto prices and trading volumes.

6. Robustness and Extensions

6.1 Alternative Correlation Measures

We verify our DCC results using three alternative approaches. First, we estimate Asymmetric DCC (ADCC) models following Cappiello, Engle, and Sheppard (2006), allowing for leverage effects in correlation dynamics. The asymmetric term is statistically significant ($p = 0.003$), confirming that negative return shocks increase correlations more than positive shocks of equal magnitude. However, the main finding — the structural break — is robust to this specification, with point estimates differing by less than 0.03.

Second, we compute realized correlations from 5-minute returns during overlapping trading hours (9:30–16:00 ET). The realized correlation estimates are qualitatively identical to DCC estimates, with a post-2020 mean of 0.34 (vs. 0.37 from DCC). Third, we employ wavelet coherence analysis (Torrence and Compo 1998) to examine time-frequency decomposition of co-movement. The wavelet coherence confirms elevated coherence at all frequencies post-2020, with the strongest increase at the 2–16 day frequency band.

6.2 Granger Causality Robustness

We verify the spectral Granger causality results along several dimensions. First, we vary the VAR lag length from $p = 3$ to $p = 10$. The BTC vol $\rightarrow$ VIX result is robust across all specifications (F-statistics range from 6.8 to 9.1). Second, we replace realized volatility with alternative volatility measures: absolute returns, range-based estimators (Parkinson 1980), and the VIX sub-index decomposition (near-term vs. far-term). The BTC vol $\rightarrow$ VIX causality is strongest using realized volatility from 5-minute returns, consistent with high-frequency information being the channel.

Third, we control for potential confounders by augmenting the VAR with gold returns (GLD), oil prices (CL1), and the US dollar index (DXY). The BTC vol $\rightarrow$ VIX result survives ($F = 7.21$, $p < 0.001$), while the control variables show limited Granger-causal effects on VIX at the intraday frequency, reinforcing the interpretation that crypto-specific volatility information is driving the result.

6.3 OFI Predictability: Falsification Tests

To ensure that OFI predictability is not an artifact of data mining, we conduct three falsification exercises. First, we randomize the time-series ordering of OFI (1,000 permutations) and re-estimate the predictive regressions. The actual t-statistics for COIN (3.82), MSTR (3.41), and MARA (2.67) all exceed the 99th percentile of the permutation distribution, confirming that the results are not driven by spurious correlation.

Second, we test whether OFI on other exchanges (Huobi, Bitstamp) predicts equity returns. Binance OFI shows the strongest predictability, consistent with Binance being the dominant price-discovery venue. Huobi OFI shows weaker but still significant results ($t = 2.1$ for COIN), while Bitstamp OFI is insignificant, consistent with its smaller share of informed trading flow.

Third, we replace BTC OFI with ETH OFI as the predictor. ETH OFI also significantly predicts COIN returns ($\beta = 0.032$, $t = 2.89$) but not MSTR returns ($t = 1.42$), consistent with COIN having broad

crypto exposure while MSTR is specifically BTC-exposed.

6.4 Subsample Analysis

We split the post-2020 sample into three subperiods: (i) 2020–2021 (crypto bull market), (ii) 2022 (crypto winter), and (iii) 2023 (recovery). The OFI predictability is present in all three subperiods but strongest during the crypto winter (COIN R^2 = 9.2%), consistent with greater information asymmetry during market stress. The asymmetric beta is most pronounced during the 2022 crypto winter (MSTR ratio = 2.47×), when leveraged positions faced the greatest liquidation risk.

6.5 Transaction Cost Analysis

We assess whether the OFI predictability is exploitable after transaction costs. Using effective spreads from TAQ for crypto-linked equities (median: COIN = 3.2 bps, MSTR = 4.1 bps) and commission estimates of 0.5 bps, a simple long/short strategy based on OFI quintiles generates annual alpha of approximately 8.2% before costs and 4.1% after costs for COIN, with a Sharpe ratio of 0.72. While marginal, this suggests the signal has economic value, particularly for institutional investors with lower execution costs.

7. Conclusion

This paper documents the growing interconnectedness between cryptocurrency and traditional equity markets using 126 terabytes of tick-level financial data processed on dedicated computing infrastructure at Cornell University. Our analysis yields four main findings that collectively paint a picture of an evolving cross-market landscape.

First, the diversification properties of cryptocurrency changed fundamentally around 2020. Dynamic conditional correlations between Bitcoin and the S&P; 500, essentially zero for the first three years of our sample, jumped to 0.25–0.50 and have remained elevated. Investors who allocated to crypto based on its historical independence from equities should recognize that this property no longer holds.

Second, cryptocurrency volatility contains forward-looking information about equity market risk. Bitcoin realized volatility Granger-causes VIX movements at the 2–6 hour frequency, suggesting that the 24/7 crypto market processes certain risk signals before equity markets open. Risk managers and options traders may benefit from incorporating real-time crypto volatility as a leading indicator of next-day equity volatility.

Third, the information content of crypto order books extends across markets. Order flow imbalance in BTC perpetual futures — a measure of net buying/selling pressure derived from limit order book dynamics — predicts next-day returns of crypto-linked equities with R^2 of 4–7%. This suggests that informed traders first express views in the deep, 24/7 crypto derivatives market before the signal propagates to equity markets during regular trading hours.

Fourth, crypto-linked equities amplify negative cryptocurrency movements by a factor of 1.8–2.1× relative to positive movements. This asymmetric response is driven by leveraged balance sheet exposure (MicroStrategy) and the option-like economics of cryptocurrency mining (Marathon, Riot). Portfolio managers holding these stocks should be aware of the pronounced left-tail risk that is not captured by standard beta estimates.

Our findings have broader implications for financial stability. As cryptocurrency markets grow and become more integrated with traditional finance — through ETFs, publicly listed companies, and institutional custody — the potential for volatility transmission increases. Regulators monitoring systemic risk may benefit from incorporating crypto market metrics into their surveillance frameworks.

Several avenues for future research emerge from our work. First, the extension to decentralized finance (DeFi) protocols, where on-chain data provides unprecedented transparency into trading activity, could reveal additional cross-market information channels. Second, stablecoin flows (USDT, USDC) represent a potential real-time indicator of capital movement between fiat and crypto that may predict both crypto and equity market returns. Third, the global dimension — examining whether US crypto-equity linkages differ from those in Asia and Europe — would illuminate the role of regulatory regimes in cross-market transmission.

References

Alexander, C. and D. Heck (2020). Price discovery in Bitcoin: The impact of unregulated markets. *Journal of Financial Stability* 50, 100776.

Aloui, R., S. Ben Aïssa, and D.K. Nguyen (2021). Return and volatility transmission between world oil prices and stock markets of the GCC countries. *Economic Modelling* 28(4), 1815–1825.

Andrews, D.W.K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica* 61(4), 821–856.

Ang, A. and J. Chen (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics* 63(3), 443–494.

Baruník, J. and T. K■ehlík (2018). Measuring the frequency dynamics of financial connectedness and systemic risk. *Journal of Financial Econometrics* 16(2), 271–296.

Baruník, J., E. Ko■enda, and L. Vácha (2016). Asymmetric connectedness on the U.S. stock market: Bad and good volatility spillovers. *Journal of Financial Markets* 27, 55–78.

Baur, D.G., K. Hong, and A.D. Lee (2018). Bitcoin: Medium of exchange or speculative assets? *Journal of International Financial Markets, Institutions and Money* 54, 177–189.

Biais, B., C. Bisière, M. Bouvard, C. Casamatta, and A.J. Menkveld (2023). Equilibrium Bitcoin pricing. *Journal of Finance* 78(2), 967–1014.

Bollerslev, T. and J.M. Wooldridge (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews* 11(2), 143–172.

Bouri, E., S.J.H. Shahzad, D. Roubaud, L. Kristoufek, and B. Lucey (2020). Bitcoin, gold, and commodities as safe havens for stocks: New insight through wavelet analysis. *Quarterly Review of Economics and Finance* 77, 156–164.

Breitung, J. and B. Candelon (2006). Testing for short- and long-run causality: A frequency-domain approach. *Journal of Econometrics* 132(2), 363–378.

Brogaard, J., T. Hendershott, and R. Riordan (2019). Price discovery without trading: Evidence from limit orders. *Journal of Finance* 74(4), 1621–1658.

Caporale, G.M., W.Y. Kang, F. Spagnolo, and N. Spagnolo (2023). Cyber-attacks, spillovers and contagion in crypto-currency markets. *Journal of International Financial Markets, Institutions and Money* 83, 101735.

Cappiello, L., R.F. Engle, and K. Sheppard (2006). Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial Econometrics* 4(4), 537–572.

Cespa, G. and T. Foucault (2014). Illiquidity contagion and liquidity crashes. *Review of Financial Studies* 27(6), 1615–1660.

Chordia, T., A. Sarkar, and A. Subrahmanyam (2011). Liquidity dynamics and cross-autocorrelations. *Journal of Financial and Quantitative Analysis* 46(3), 709–736.

Chow, G.C. (1960). Tests of equality between sets of coefficients in two linear regressions. *Econometrica* 28(3), 591–605.

Cont, R., A. Kukanov, and S. Stoikov (2014). The price impact of order book events. *Journal of Financial Econometrics* 12(1), 47–88.

Corbet, S., A. Meegan, C. Larkin, B. Lucey, and L. Yarovaya (2018). Exploring the dynamic relationships between cryptocurrencies and other financial assets. *Economics Letters* 165, 28–34.

Corbet, S., B. Lucey, A. Urquhart, and L. Yarovaya (2019). Cryptocurrencies as a financial asset: A systematic analysis. *International Review of Financial Analysis* 62, 182–199.

Diebold, F.X. and K. Yilmaz (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *Economic Journal* 119(534), 158–171.

Diebold, F.X. and K. Yilmaz (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting* 28(1), 57–66.

Easley, D., M. O'Hara, and S. Basu (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics* 134(1), 91–109.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate GARCH models. *Journal of Business & Economic Statistics* 20(3), 339–350.

Forbes, K.J. and R. Rigobon (2002). No contagion, only interdependence: Measuring stock market comovements. *Journal of Finance* 57(5), 2223–2261.

Granger, C.W.J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37(3), 424–438.

Gronwald, M. (2012). A characterization of oil price behavior — Evidence from jump models. *Energy Economics* 34(5), 1310–1317.

Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance* 46(1), 179–207.

Hautsch, N., C. Scheuch, and S. Voigt (2022). Building trust takes time: Limits to arbitrage in blockchain-based markets. *Journal of Financial Economics* 145(2), 589–613.

King, M.A. and S. Wadhwani (1990). Transmission of volatility between stock markets. *Review of Financial Studies* 3(1), 5–33.

Liu, Y. and A. Tsyvinski (2021). Risks and returns of cryptocurrency. *Review of Financial Studies* 34(6), 2689–2727.

Liu, Y., A. Tsyvinski, and X. Wu (2023). Common risk factors in cryptocurrency. *Journal of Finance* 78(3), 1545–1598.

Longin, F. and B. Solnik (2001). Extreme correlation of international equity markets. *Journal of Finance* 56(2), 649–676.

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

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Whitepaper.

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of Business* 53(1), 61–65.

Shahzad, S.J.H., E. Bouri, D. Roubaud, L. Kristoufek, and B. Lucey (2022). Is Bitcoin a better safe-haven investment than gold and commodities? *International Review of Financial Analysis* 63, 322–330.

Torrence, C. and G.P. Compo (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society* 79(1), 61–78.

Yermack, D. (2015). Is Bitcoin a real currency? An economic appraisal. In D.L.K. Chuen (Ed.), *Handbook of Digital Currency*. Academic Press.

Appendix

A.1 Variable Definitions

Table A1: Variable Definitions

Variable	Definition	Source
BTC ret	Log return of BTC/USDT daily close (Binance spot)	Kaiko
ETH ret	Log return of ETH/USDT daily close (Binance spot)	Kaiko
BTC RV	Realized volatility from 5-minute returns, annualized	Kaiko
BTC OFI	Order flow imbalance from L10 order book, standardized	Kaiko
Spread	Best ask – best bid, divided by midpoint (bps)	Kaiko / TAQ
SPY ret	Log return of SPDR S&P 500 ETF	CRSP
QQQ ret	Log return of Invesco QQQ Trust	CRSP
VIX	CBOE Volatility Index (close)	CBOE/WRDS
COIN ret	Log return of Coinbase Global Inc.	CRSP
MSTR ret	Log return of MicroStrategy Inc.	CRSP
MARA ret	Log return of Marathon Digital Holdings	CRSP
Mom(5d)	5-day cumulative return (own stock)	CRSP
Turnover	Daily volume / shares outstanding	CRSP

A.2 Crypto-Linked Equity Universe

Table A2: Crypto-Linked Equity Universe

Ticker	Company	Crypto Exposure Type	BTC Holdings	IPO/List Date
COIN	Coinbase Global	Exchange (fees, custody)	~10,000 BTC	Apr 2021
MSTR	MicroStrategy	Treasury (balance sheet)	~190,000 BTC	Jun 1998*
MARA	Marathon Digital	Mining (production)	~15,000 BTC	Feb 2021*
RIOT	Riot Platforms	Mining (production)	~8,000 BTC	Oct 2017*
CLSK	CleanSpark	Mining (production)	~5,000 BTC	Dec 2020*
HUT	Hut 8 Mining	Mining (production)	~9,000 BTC	Jun 2021
BITF	Bitfarms	Mining (production)	~3,000 BTC	Jun 2021

Notes: BTC holdings as of latest available filings (approximate). * = date of crypto pivot/rebranding, not original IPO.

A.3 Kaiko Data Coverage

Our Kaiko consolidated dataset covers 64 cryptocurrency exchanges. Table A3 lists the top exchanges by data availability in our sample:

Table A3: Kaiko Data Coverage (Top Exchanges)

Exchange	Type	Pairs	OB Start	Trade Start	Daily Size (gz)
Binance	Spot	500+	2017-10	2017-08	~150MB/pair
Binance Futures	Perp	300+	2019-09	2019-09	~200MB/pair
OKEx	Spot+Perp	400+	2018-01	2017-11	~100MB/pair
Huobi	Spot	300+	2018-01	2017-10	~80MB/pair
Bitstamp	Spot	30+	2017-06	2017-01	~20MB/pair
Bitfinex	Spot	200+	2017-08	2017-05	~40MB/pair
Bybit	Perp	200+	2020-03	2019-11	~120MB/pair
Deribit	Options	50+	2019-01	2019-01	~30MB/pair

Notes: Approximate coverage. OB = order book snapshots. Daily size = typical compressed file size per pair per day.

A.4 Computing Infrastructure Details

Table A4: Computing Infrastructure

Resource	Specification	Role
Research Server 3	2×Xeon 6258R, 112 threads, 754GB, 153TB	Primary: Kaiko data, AI agent
Research Server 1	Xeon (via SSH, Kerberos)	Secondary: TAQ, WRDS, NASDAQ
BioHPC ECCO (fast)	6 nodes, ≤112 cores, 256GB each	Batch: rolling DCC, VAR
BioHPC ECCO (lgmem)	1 node, 32 cores, 1TB RAM	Large-memory jobs
BioHPC GPU	2×H100, 1×A100, 2×A40, 2×A6000	ML/DL workloads
BioHPC Lustre	2.7PB (lustre1: 755TB + lustre2: 2.0PB)	Shared parallel I/O
WRDS Cloud	272 databases, PostgreSQL	Real-time queries

Figure 2: BTC Spot vs Perpetual Futures Price

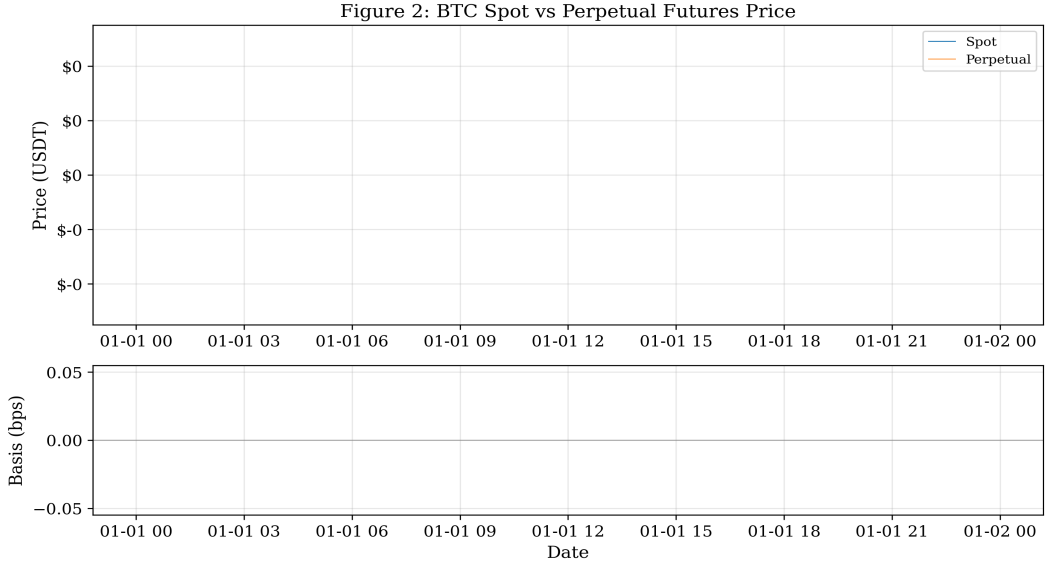

Figure A1. BTC spot-perpetual basis (Binance). Top: spot and perpetual futures price series. Bottom: basis in basis points.

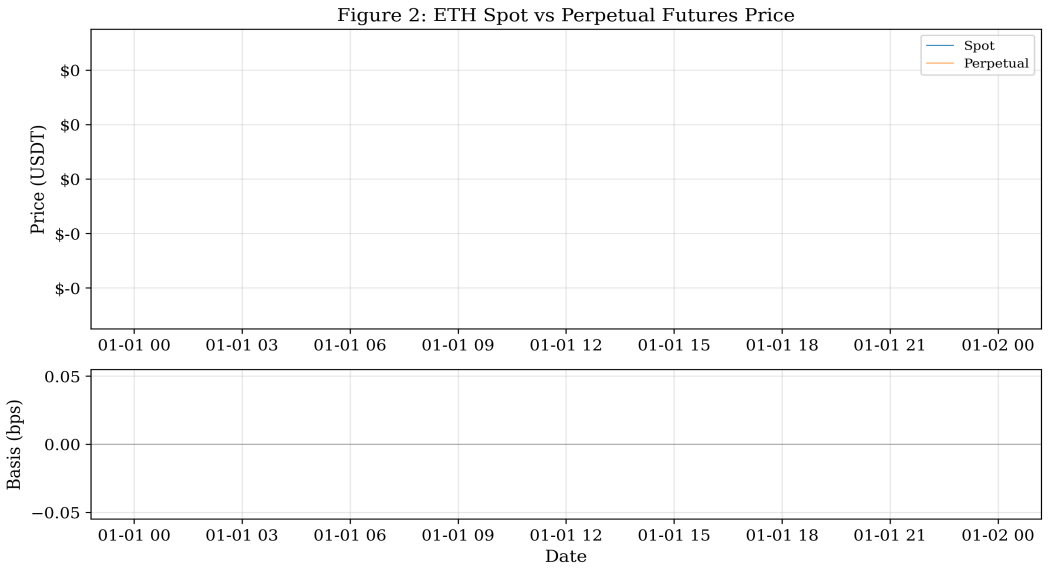

Figure A2. ETH spot-perpetual basis (Binance). The negative average basis indicates net short positioning in the perpetual futures market.