# New Limits to Arbitrage: Evidence from Crypto Perpetual Futures Markets.

Don Chance* Ravi Joshi†

December 4, 2025

**Abstract**

Perpetual futures rely on funding rates for their arbitrage mechanism. Using Terra, FTX, and Silicon Valley Bank collapses as natural experiments and crypto perpetual futures as a laboratory, we document an anomaly: similarly severe crises produce opposite effects on arbitrage functionality in perpetual futures. Applying the global games framework, we model arbitrageurs' stay-or-exit decisions as functions of a public signal (spot volatility) and private signals (adverse selection costs). We identify a volatility threshold at which the marginal effect of adverse selection on arbitrage capital flips from stabilizing to destabilizing—triggering a coordination-driven run. During Terra's collapse, Ethereum's volatility exceeded this threshold (1.8%/hour); open interest fell by 65%, and funding rates ceased correcting the basis. During the FTX and SVB episodes, we find no significant threshold despite comparable severity; arbitrage effectiveness remained intact and funding rates corrected basis. Our findings demonstrate that arbitrage breakdowns arise from endogenous coordination thresholds during periods of market stress, bridging classical limits-to-arbitrage theory with global games coordination frameworks.

*dchance@lsu.edu
†rjosh15@lsu.edu

# 1 Introduction

Perpetual futures (perps), constituting approximately 93% of the crypto derivative's trading volume, with daily flows often exceeding $100 billion[1] , are a type of derivative contract that enables traders to speculate on the price of an asset—like Bitcoin, Ethereum, or other cryptocurrencies—without needing to buy or own the underlying asset itself. Unlike traditional futures contracts, perpetual futures don't have an expiration date, meaning traders can hold their positions indefinitely (Shiller, 1993), as long as they maintain the required margin.Because of no expiration date, there is no guaranteed date when the futures price would converge to the spot price. The fundamental mechanism anchoring perpetual futures to spot markets is the funding rate [2]—a periodic cash flow between long and short positions that incentivizes arbitrageurs when the perpetual deviates from the underlying (Angeris et al., 2023). When futures trade above spot, long positions periodically fund shorts, and vice versa, thereby closing the price deviations between perp and spot. In essence, the funding rate is the primary arbitrage mechanism in perpetuals.

In well-functioning markets, funding rates effectively enforce price parity (He et al. (2022)). However, during periods of market stress, the mechanism may fail. Using the global games framework (Goldstein and Pauzner, 2005; Morris and Shin, 2001) and using Ethereum perpetuals as a laboratory, we test the determinants of arbitrage limits in the perpetual futures market during periods of market stress by examining open interest, which captures arbitrage capital available. Open interest captures the total number of outstanding contracts that remain unsettled, and is a standard indicator for market participation, liquidity, leverage, and arbitrage capital in derivatives markets. We examine the marginal effect of information costs, and

[1] https://business.cornell.edu/article/2025/02/perpetual-futures-contracts-and-cryptocurrency/

[2] https://www.binance.com/en/support/faq/detail/360033525031

the volatility of the underlying on open interest. The marginal effect captures the response of arbitrage capital to a marginal change in adverse-selection cost, given the current volatility level. We show that there exists a volatility threshold beyond which adverse selection becomes capital-destructive, and even short breaches of that threshold trigger persistent deleveraging, generating a run-like collapse in open interest. This creates capital limits to arbitrage (Shleifer and Vishny, 1997), and funding rates stop responding to basis. In other words, capital constraints on arbitrage driven by spot volatility thresholds weaken arbitrage mechanism.

We exploit the most dramatic collapses of cryptocurrencies that serve as natural experiments to study limits to arbitrage in perpetual futures: the collapse of stablecoin Terra and the collapse of crypto exchange FTX. We also study the Silicon Valley Bank (SVB) collapse, as SVB had significant exposure to crypto-related firms and deposits—most notably Circle's USDC reserves—making its failure tightly linked to cryptocurrency market stress. The duration of stress episodes was very similar, and all destroyed similar market value within a week ($60B Terra, $32B FTX, $42B SVB). Stablecoin Terra and crypto exchange FTX failed within three days, and the run on Silicon Valley Bank (SVB) was a two-day event. We first show that the funding rate remains uniform in pre-crash periods but becomes volatile and drops sharply once the crash begins. The results are consistent across all stress episodes examined in our study. This suggests that the arbitrage mechanism responsible for closing the basis weakens once the stress period starts. Post-crash, the funding rates revert to pre-crash levels, suggesting that the arbitrage mechanism begins to function well enough to close the basis.

Next, we examine funding rate elasticity during different market regimes. Funding-rate elasticity captures how closely the funding rate adjusts to the basis, and is computed as the coefficient from regressing funding rates on the basis. A theoretical elasticity of one implies that the arbitrage mechanism fully aligns the perpetual price with the spot price and effectively closes the basis. In essence,

funding-rate elasticity is a measure of arbitrage mechanism effectiveness. We find that the elasticity is approximately 0.64 before the Terra shock, collapses to 0.38 during the May 7–9 crash, and rebounds to 0.91 afterward. This pattern confirms that the arbitrage effectiveness nearly vanished during Terra's crash but quickly recovered post-crash. However, the value of dynamic funding elasticity was consistently over the theoretical value of one during the periods of the FTX and SVB crashes. Pre FTX and SVB crashes, we see a stable and near-proportional relation between funding rate and the basis. During the crash period, the funding mechanism remained highly elastic but was subject to a level shift— funding rates increased (futures were trading above(SVB), below(FTX) spot), and the arbitrage mechanism was closing the basis. Post-crash, funding rates became negative (Terra and SVB), but funding elasticity became stable, consistent with a restoration of equilibrium. Our central finding is an anomaly: three comparably severe market stress events—Terra on one hand, and FTX and SVB on the other—produce opposite effects on funding rate elasticity in Ethereum perpetual futures. Arbitrage effectiveness collapses only during the period of the Terra crash, but remains intact during the FTX and SVB crash episodes.

To understand the cause of the anomaly, we first examine open interest surrounding the crash episodes. Open interest, in the derivatives market, reflects the depth of speculative and arbitrage capital available to absorb price and funding-rate adjustments. When open interest declines, the effective arbitrage capacity of the market contracts. We find that the three crisis events exhibit starkly different open interest dynamics that align precisely with their differential arbitrage effectiveness patterns. During the Terra collapse, open interest experienced a catastrophic collapse. From a stable pre-crash level averaging $58.7 million in notional value, open interest briefly spiked to $66 million before plummeting to just $21 million—a 65% decline that persisted throughout the post-crash period with minimal recovery. This wholesale exodus of capital represents not merely a position reduction but a

fundamental breakdown in market-making and arbitrage infrastructure. In sharp contrast, the FTX bankruptcy and SVB failure displayed remarkable resilience in open interest despite their comparable or even greater systemic importance. During the FTX crisis, open interest increased only marginally from $9.5 million to over $13 million, and exceeded pre-crisis levels. The SVB episode showed even greater stability, with open interest fluctuating by just $0.75 million during the crash period. These differential patterns in market participation directly map to our funding-rate elasticity findings: arbitrage effectiveness collapse during Terra's episode coincides with the flight of arbitrage capital, while maintained effectiveness during FTX and SVB reflects the continued arbitrage capital to intermediate basis divergences.

To examine the cause of the collapse in the open interest rate, we utilize the global games approach (Goldstein and Pauzner, 2005; Morris and Shin, 2001) by examining the impact of information costs and market forces on open interest, the arbitrage capital. Information asymmetry creates strategic complementarities among arbitrageurs: when adverse selection costs are high, each arbitrageur rationally anticipates that others may possess superior information about fundamental value shifts, leading to coordinated withdrawal of arbitrage capital((Brunnermeier and Pedersen, 2005)). In our empirical estimation, the fundamental is the resilience of the spot market, which we proxy with short-horizon spot volatility $\sigma_t$ that acts as a public signal, so that increases in $\sigma_t$ are interpreted as a deterioration in underlying market fundamentals. Adverse selection costs reflect a private signal about the informativeness and toxicity of spot trades. Open interest is the aggregate outcome of individual stay-or-exit decisions. The interaction between volatility and adverse selection costs in our model, therefore, provides an empirical analogue of the global-games coordination mechanism.

Our estimates imply a volatility threshold $\sigma^*$ such that, when volatility is below $\sigma^*$, marginal increases in adverse selection costs are absorbed by arbitrage capital. Once volatility exceeds $\sigma^*$, however, the marginal effect of higher adverse

selection costs turns negative: additional informational frictions are associated with a sharp decline in open interest. In this high-volatility region, small increases in adverse selection costs coincide with a coordinated withdrawal of arbitrage capital, consistent with a global-games–style run equilibrium.

We find that adverse selection costs exhibit dramatically different patterns across the three crisis episodes. During the Terra collapse, median adverse selection costs surged from \$1,434 to \$12,911 per hour, a ninefold increase. The FTX bankruptcy generated even higher absolute adverse selection costs (median: \$14,562 during crash). While the median exceeded Terra's, the persistence and uniformity differed markedly—costs rapidly normalized post-crisis to \$980, below pre-crisis levels of \$1,989. This V-shaped recovery indicates that once the FTX fraud was revealed, information asymmetry actually decreased relative to the pre-crisis period as uncertainty resolved into common knowledge. The SVB failure presents a striking contrast: despite triggering a regional banking crisis, adverse selection costs barely moved, declining slightly from \$486 to \$456 during the crisis phase.

If volatility or other market forces were driving open interest collapse, these variables should load strongly and consistently on open interest across all events. Instead, we uncover a distinctive pattern during the Terra crash, adverse selection costs become significant only conditional on volatility, with volatility ($\sigma^*$) threshold of approximately $\approx 1.8\%$ per hour at which adverse selection costs do not impact open interest. Below this volatility level, adverse-selection costs are associated with higher open interest. But once volatility spikes past $\sigma^*$, the sign flips, and adverse selection costs put negative pressure on open interest, consistent with deleveraging that we find. In contrast, FTX and SVB episodes do not exhibit a statistically or economically meaningful $\sigma^*$, suggesting that not all crashes cross into the run region consistent with the global-games framework (Goldstein and Pauzner, 2005).

The robustness of this coordination mechanism is further validated through an out-of-sample test using the 2021 China regulatory crackdown, an event that introduced sudden, policy-driven information asymmetry. Consistent with the findings during Terra collapse episode, the crackdown's initial phase (Phase I) saw a sharp collapse in funding elasticity and a significant 45% decline in open interest, accompanied by the robust identification of a statistically significant coordination threshold ($\sigma^* \approx 4.6\%$/hour) which was breached by realized volatility only during this phase of high information asymmetry. Crucially, the $\sigma^*$ threshold was statistically insignificant during the subsequent phases of the China crackdown (Phases II, III, and IV), mirroring the results of the FTX and SVB periods where the coordination-driven run was not triggered and open interest remained stable. This replication across an independent, policy-driven shock confirms that capital constraints to arbitrage are dependent on this endogenous coordination threshold, providing powerful support for limits to arbitrage (Shleifer and Vishny, 1997).

In the final part of the paper, we examine the consequences of these runs for the funding mechanism itself. We show that when open interest collapses, the pass-through from basis to funding rate—our measure of arbitrage effectiveness—declines sharply. In other words, runs on arbitrage capital not only shrink positions but also weaken the capacity of funding rates to realign futures prices with spot fundamentals. This behavior is reminiscent of run equilibria and should be relevant for equity, FX, and commodity futures markets during stress episodes. This connects the global-games run mechanism directly to the functioning of the core pricing anchor mechanism in derivatives.

The rest of the paper is organized as follows: Section 2 provides the literature review. Section 3 presents data and summary statistics. Section 4 presents empirical methodology and results. Section 5 presents results. Section 6 presents robustness tests. Section 7 concludes.

## 2 Literature Review

The arbitrage mechanism anchoring perpetual futures to spot markets is crucial for market efficiency, yet it is inherently vulnerable to the classic "limits to arbitrage" documented in financial theory. The earliest work by (Grossman and Stiglitz, 1980) established that fully efficient markets are impossible if information acquisition is costly, creating a fundamental limit on how closely prices can track fundamentals. More formally, (Shleifer and Vishny, 1997) showed that capital constraints on arbitrageurs, such as funding constraints, can cause asset prices to deviate persistently from fundamental value. The idea of perpetual futures was formalized in academic literature by (Shiller, 1993), but the practical application of (Shiller, 1993) concept lay dormant until cryptocurrency exchanges operationalized perpetual futures, sparking a new wave of research in this direction. While recent research on perpetuals (Ackerer et al., 2024; He et al., 2022; Gornall et al., 2024) establish the general effectiveness of the arbitrage mechanism, they leave open the question of how and why it fails during systemic crises, a gap this paper aims to fill.

The central mechanism investigated in this paper is the endogenous failure of arbitrage driven by coordination runs, a concept pioneered in the literature on bank and debt runs. This framework, formalized by (Morris and Shin, 2001) and applied to banking by (Goldstein and Pauzner, 2005), models agents' decisions (to stay or run) as strategic complements, where individual action is dependent on the perceived actions of others. In these Global Games (GG) models, a unique coordination threshold ($\sigma^*$) exists, based on the interaction between a public signal (market stress) and private signals (information costs). When the public signal crosses this threshold, strategic complementarities flip, leading to a coordinated and destructive withdrawal of capital. This paper extends this framework to model the arbitrageurs' 'stay-or-exit' decision in the perpetual futures market, by providing reduced-form model to capture the public signal of the run threshold.

In conventional asset classes (FX, commodities, etc.), perpetual futures have not been traded in practice, but analogous principles exist. The closest parallel is the forward/futures pricing parity enforced by arbitrage. For example, in FX markets, Covered Interest Parity (CIP) ensures that the forward exchange rate is tied to the interest rate differential between currencies – any deviation is quickly arbitraged away under normal conditions. (Du et al., 2018) document how CIP famously failed during the 2008 crisis, with forward prices diverging from parity by dozens of basis points as credit risk and funding constraints impeded arbitrage. We extend this literature by examining the same mechanism in the crypto derivative market. (De Blasis and Webb, 2022) compare the contract design and microstructure of Bitcoin quarterly futures versus perpetual futures. They document episodes where arbitrageurs could exploit mispricings between the two markets, but as crypto markets matured, those opportunities became less frequent. (Kim and Park, 2025) provide a theoretical framework for setting the funding rate process optimally and discusses path-dependent funding and hedging for exchanges to ensure the perp price stays aligned with the underlying.

(Streltsov and Ruan, 2022) focus on how the introduction of perpetuals changes market behavior. Using natural experiments, they found that when exchanges added perpetual futures, trading activity surged and those markets attracted more informed traders. (Alexander et al., 2020) examine BitMEX Bitcoin futures (both the perpetual swap and fixed-maturity quarterly futures) and find that the BitMEX perpetual plays a dominant role in price discovery for Bitcoin, often leading the spot market. Their results also suggest the periods of information inefficiency, where futures and spot diverge, profitable hedging and arbitrage opportunities are created. Broadly, several papers have explored the role of arbitrage in digital assets Liu et al. (2023); Joshi (2025); Makarov and Schoar (2020); Lyons and Viswanath-Natraj (2023). We extend the literature by examining the arbitrage mechanism in perpetual futures during periods of systemic stress. Most importantly, we extend

literature on classical limits to arbitrage (Shleifer and Vishny, 1997; Grossman and Stiglitz, 1980) by showing that arbitrage capital constraints can be an endogenous result of coordination failure, and we provide an empirical application of Global Games coordination thresholds (Goldstein and Pauzner, 2005; Morris and Shin, 2001) to the perpetual futures market.

# 3 Data & Sample Construction

We utilize minute-by-minute trade and perpetual futures data for the Ethereum/USDC pair from Kaiko, a leading provider of cryptocurrency data (Makarov and Schoar, 2020). The data spans 2020 through 2023, encompassing multiple stress episodes. We use data for the Kraken exchange because its hourly funding (AsiaNext Exchange, 2024) provides us with a setup to examine high-frequency data. We compute market microstructure variables from ETH–USDC trade data and aggregate them hourly. Following Hasbrouck (1993), the realized volatility of Ethereum prices, $\sigma_{ETH,t}$, is obtained using a 24-hour rolling window of hourly close-to-close price changes. This estimator adjusts for the bias induced by bid-ask bounce and provides an estimate of the efficient price volatility. This approach builds on the model of Roll (1984) but directly estimates the volatility of the efficient price rather than the spread. The variable is measured in USD and is computed as:

$$\sigma_{ETH,t} = \sqrt{\mathrm{Var}(\Delta P_t) + 2, \mathrm{Cov}(\Delta P_t, \Delta P_{t-1})}, \qquad \Delta P_t = \ln\left(\frac{\mathrm{Close}_t}{\mathrm{Close}_{t-1}}\right). \tag{1}$$

We compute hourly momentum, $Mom_t$, as the logarithmic price change from the previous hour, capturing short-term return persistence.

$$Mom_t = \ln\left(\frac{\mathrm{Close}_t}{\mathrm{Close}_{t-1}}\right), \tag{2}$$

Order flow imbalance, $OF_{i,t}$, measures net buying pressure as the signed volume imbalance within each hour,

$$OF_{i,t} = \frac{\sum_{j \in t} q_j \, V_j}{\sum_{j \in t} |V_j|}, \tag{3}$$

where $q_j \in \{+1, -1\}$ indicates buyer- or seller-initiated trades and $V_j$ is the trade volume.

## 3.1 Estimation of Adverse Selection

The measure of information asymmetry follows (Glosten and Harris, 1988) framework to compute adverse selection, which estimates the adverse-selection component of the bid–ask spread by regressing price changes on trade direction and trade size. Accordingly, we denote the unobserved true price of the asset by $m_t$. It represents the price of the asset assuming a fully competitive market maker and no inventory costs or clearing fees. Innovations in $m_t$, therefore, result from public news or information arrival through order flows. More formally,

$$m_t - m_{t-1} = e_t + Q_t Z_t, \tag{4}$$

with $Q_t = 1(-1)$ for a buyer- (seller-) initiated trade. In Equation 4, $e_t$ represents the impact of public news, and $Z_t$ is the market maker's compensation for bearing the adverse selection risk. Hence, $Z_t$ is the adverse selection component of the bid-ask spread. The observed prices reflect factors such as the market maker's monopoly power, inventory costs, and clearing fees. To account for these factors, Glosten and Harris (1988) include a second component in their specification, called

the *transitory component*. Formally,

$$P_t = m_t + Q_t C_t, \tag{5}$$

where the dependence of the transitory component on the trade direction reflects the fact that market makers buy low and sell high.

Larger trades increase the bid-ask spread through both components. Therefore, we allow $Z_t$ and $C_t$ to depend on the volume, $V_t$. Glosten and Harris (1988) assume a linear dependence to facilitate the estimation.

$$Z_t = z_0 + z_1 V_t$$
$$C_t = c_0 + c_1 V_t$$

Substituting these into Equation 5 and taking the first difference gives

$$P_t - P_{t-1} = c_0(Q_t - Q_{t-1}) + c_1(Q_t V_t - Q_{t-1} V_{t-1}) + z_0 Q_t + z_1 Q_t V_t + e_t. \tag{6}$$

We estimate the parameters in Equation 6 for every trading hour. We observe all the variables in Equation 6 in our data and can estimate the parameters, $c_0, c_1, z_0$, and $z_1$ through OLS regressions.[3] Using the estimated parameters, $\hat{c_0}, \hat{c_1}, \hat{z_0}$, and $\hat{z_1}$, we calculate $\hat{C}_t$ and $\hat{Z}_t$ for each trade. Our measure of the adverse selection component of the bid-ask spread for the trading hour $h$ is

$$Adverse\ Selection_h = \sum_{t \in h} \hat{Z}_t$$

[3]Contrary to Glosten and Harris (1988), we can observe the trade direction ($Q_t$) in my data. Furthermore, the rounding error in prices is negligible in my setup, given the minuscule tick size on Exchanges (0.00000010 for BTC/USD pair). Therefore, we do not resort to the maximum likelihood method to estimate the parameters.

All variables are computed hourly and lagged one period when used as predictors in the model.

## 3.2 Summary Statistics

Table 1 reports summary statistics for Ethereum spot and perpetual markets across three major disruption events: the Terra/Luna collapse (May 2022), the FTX bankruptcy (November 2022), and the Silicon Valley Bank (SVB) failure (March 2023). Each event is partitioned into pre-crisis, crisis, and post-crisis windows, allowing for a cross-event comparison of market microstructure dynamics under extreme stress. All statistics are based on hourly observations. Adverse selection costs are reported in scaled units of cost per \$10,000 notional per hour; specifically, the table entries must be multiplied by $10^4$ to obtain the underlying cost. For example, an entry of 1.69 corresponds to an underlying adverse-selection cost of $1.69 \times 10^4 = 16{,}900$ USD per hour.

Several patterns emerge. First, spot prices decline sharply in the crisis windows for all three events. Mean Ethereum prices fall by about 23% during Terra (from \$2,913 to \$2,256), 5% during FTX (from \$1,442 to \$1,377), and 9% during SVB (from \$1,630 to \$1,485). Post-crisis, prices remain below their pre-crisis levels in all three cases, with average discounts of roughly 33% for Terra, 14% for FTX, and 6% for SVB. These persistent valuation gaps suggest that major crypto disruptions are associated with lasting downward price adjustments rather than rapid mean reversion.

Funding-rate and basis dynamics reveal distinct patterns in market positioning across events. In Terra, the mean funding rate rises from approximately 0.002 in the pre-crisis window to about 0.01 during the crash, before turning negative in the post-crisis period ($-0.002$). Combined with a shift in the average basis from mildly positive (0.17) to more elevated levels (0.32) during the crash, then slightly

negative ($-0.05$) afterward, this pattern is consistent with an initial buildup of leveraged long positions supporting arbitrage, followed by broad deleveraging once the peg breakdown becomes persistent.

By contrast, the FTX episode is characterized by persistently negative funding rates and deeply negative bases. The mean funding rate falls from $-0.003$ pre-crisis to around $-0.02$ in the crash window and remains negative post-crisis. The basis moves from a modest discount ($-0.11$) to a much larger one ($-0.85$) during the crash, before settling at an intermediate discount ($-0.47$) post-crisis. This pattern reflects sustained short pressure and a prolonged period in which perpetual prices trade at a discount to the index. The SVB disruption shows a more modest but highly asymmetric dislocation: the pre-crisis basis is slightly positive (0.16), turns mildly negative in the crash window ($-0.06$), and collapses to a large discount ($-2.65$) in the post-crisis period, indicating persistent pricing stress even after the immediate failure window.

Volatility and transaction costs rise across all three events. Measured by $\sigma_{\text{ETH}}$, average volatility roughly doubles in the crash window for both Terra and FTX (from about 1% to 2%), and increases from roughly 0.5% to 1% for SVB. Post-crisis, volatility normalizes quickly for FTX, falling back toward pre-crisis levels, but remains at or above pre-crisis levels for Terra and especially SVB.

Bid–ask spreads and adverse-selection-driven trading costs (market makers' risk premia) show particularly strong stress patterns once the scaling is taken into account. For Terra, the mean risk premia rises from 2700 USD in the pre-crisis window to 18,400 USD during the crash, before falling to 3600 post-crisis. FTX displays the same qualitative pattern: Risk premia rises from 6,800 to 14,000 USD during the crash and then drops to 1,800 post-crisis. SVB exhibits a smaller but still meaningful increase in adverse selection costs from 700 pre-crisis to 1,800 USD in the crash window, followed by a modestly higher post-crisis level of 2,200. Across all events, market makers' risk premia becomes markedly more expensive.

Open interest and order flow highlight differences in market participation and the depth of arbitrage capital across events. Terra exhibits a dramatic collapse in mean open interest, from 58.74 million pre-crisis to 49.27 million during the crash and just 20.30 million post-crisis, implying a sharp contraction in the capacity of speculative and arbitrage capital to absorb price and funding-rate shocks. In contrast, FTX shows a mild increase in open interest from 10.96 to 11.33 million in the crash window and to 11.83 million post-crisis, suggesting that some capital remained engaged or was even attracted by the dislocation. SVB features a gradual erosion of open interest (from 9.36 to 9.05 and then 8.94 million), consistent with a more conventional risk-off retrenchment.

Order-flow imbalances become more negative during all three crisis windows, reflecting the dominance of aggressive selling pressure. Terra's mean order flow moves from approximately zero (0.002) pre-crisis to a clearly negative value (0.07) during the crash, before turning positive (0.08) post-crisis, suggesting net buying interest as prices stabilize at lower levels. FTX shifts from slightly positive order flow (0.06) before the collapse to negative order flow (0.04) in the crash window, and remains marginally negative thereafter, consistent with continued net selling and unwinding of long exposure after the bankruptcy. SVB displays the most severe deterioration in order flow in relative terms, declining from 0.10 pre-crisis to 0.25 in the crash period, with only a partial recovery (0.04) post-crisis. These patterns underscore that microstructure responses are highly event-specific, shaped by both the nature of the underlying shock and the resilience of arbitrage capital.

# 4 Empirical Setting

Our empirical model is motivated by global-games run frameworks in the spirit of Morris and Shin (2001); Goldstein and Pauzner (2005). In those models, agents

decide whether to *stay* or *run* based on a latent fundamental $\theta_t$ and noisy signals about that fundamental. Strategic complementarities in action give rise to a unique threshold $\theta$ at which the equilibrium switches from no-run to run. We map this structure to the perpetual market as follows:

- The *fundamental* is the resilience of the ETH spot–perpetual complex. We proxy resilience by short-horizon spot volatility $\sigma_t$. Higher $\sigma_t$ indicates a more fragile trading environment with greater execution and margin risk, and thus weaker fundamentals. Volatility $\sigma_t$ is observed by all market participants and therefore acts as a *public stress signal*.

- Each arbitrageur faces *private information costs* summarized by the adverse selection cost per trade. This measure captures market makers' risk premia from trading against better-informed counterparties. It plays the role of a *private signal* to market participants.

- Open interest $OI_t$ is the aggregate outcome of individual stay-or-exit decisions. A run on arbitrage capital corresponds to a sharp decline in $OI_t$ as many arbitrageurs simultaneously choose to exit.

In a global-games interpretation, each arbitrageur observes the public signal $\sigma_t$ and a noisy private signal about information costs, and chooses whether to maintain or liquidate their position. Under standard assumptions, the equilibrium can be characterized by a threshold rule: for a given level of public stress $\sigma_t$, arbitrageurs exit when their private information cost exceeds a cutoff. Arbitrage capital outflow, open interest decline, is a function of $\sigma_t$ and the distribution of private signals.

## 4.1 Reduced-Form Specification for Open Interest

Rather than estimating a structural global game, we work with an empirical reduced-form representation of the equilibrium relationship between open interest, market frictions, and adverse selection costs. We estimate open interest as

$$OI_t = \alpha + \beta_{AS} \cdot \text{AS\_cost}_{t-1} + \beta_\sigma \cdot \sigma_{t-1} + \beta_{AS\times\sigma} \cdot \text{AS\_cost}_{t-1} \times \sigma_{t-1} + \gamma' X_{t-1} + \mu_{h(t)} + \varepsilon_t, \quad (7)$$

where:

- $OI_t$ is open interest (in millions of USD) in the ETH perpetual at hour $t$;
- $\text{AS\_cost}_t$ is the adverse selection cost ;
- $\sigma_t$ is the spot volatility of ETH;
- $\text{AS\_cost}_t \cdot \sigma_t$ is the interaction that captures state-dependence of the effect of information costs on arbitrage capital;
- $X_t$ are controls (such as order-flow imbalance, and momentum);
- $\mu_{h(t)}$ are hour-of-day fixed effects;
- $\varepsilon_t$ is an error term.

Equation (7) can be interpreted as a local linear approximation to the global-games equilibrium mapping from the public signal $\sigma_t$ and the distribution of private information costs (summarized by $\text{AS\_cost}_t$) into the aggregate outcome $OI_t$. The key nonlinearity implied by global-games theory is that the impact of information costs on the equilibrium mass of "stayers" (arbitrageurs that keep positions) is *state-dependent*: beyond a certain level of public stress, small changes in private information costs can trigger a coordinated run.

## 4.2 Marginal Effect of Adverse Selection and Volatility Threshold

The global-games interpretation focuses on how the marginal effect of information costs on the aggregate action depends on the public signal. In our specification, the marginal effect of adverse selection costs on open interest is given by the partial derivative of $OI_t$ with respect to $\text{AS\_cost}_t$:

$$\frac{\partial OI_t}{\partial \text{AS\_cost}_t} = \beta_{AS} + \beta_{AS\times\sigma} \cdot \sigma_t. \tag{8}$$

Equation (8) has a simple interpretation:

- For a fixed level of volatility $\sigma_t$, $\frac{\partial OI_t}{\partial \text{AS\_cost}_t}$ measures the expected change in arbitrage capital (open interest) associated with a marginal increase in information costs.
- If $\frac{\partial OI_t}{\partial \text{AS\_cost}_t} > 0$, the market can absorb higher information costs without losing arbitrage capital; adverse selection is not yet "run-inducing."
- If $\frac{\partial OI_t}{\partial \text{AS\_cost}_t} < 0$, higher information costs are associated with a decline in open interest; adverse selection becomes *capital-destructive* and is consistent with a run on arbitrageurs' positions.

A central implication of global-games models is the existence of a threshold in the public signal at which behavior changes discontinuously in equilibrium. In our reduced-form representation, this threshold is captured by the volatility level at which the marginal effect of adverse selection switches sign. Equating (8) to zero yields the *threshold volatility*. The volatility threshold $\sigma^*$ in (9) is then the empirical analogue of the global-games run threshold: a level of public stress at which a marginal increase in private information costs changes from stabilizing to destabilizing for arbitrage capital.

$$\frac{\partial OI_t}{\partial \text{AS\_cost}_t} = 0 \quad \Longrightarrow \sigma^* = -\frac{\beta_{AS}}{\beta_{AS\times\sigma}} \tag{9}$$

For $\sigma^*$ to be a strong and valid threshold, $\beta_{AS}$ and $\beta_{AS\times\sigma}$ must be of opposite signs and be statistically significant—the implied $\sigma^*$ in (9) is positive. This delivers the global-games–style threshold:

$$\frac{\partial OI_t}{\partial \text{AS\_cost}_t} = \begin{cases} > 0, & \text{if } \sigma_t < \sigma^* \rightarrow No\ Run \\ < 0, & \text{if } \sigma_t > \sigma^* \rightarrow Run \end{cases} \tag{10}$$

When $\sigma_t < \sigma^*$, marginal increases in adverse selection costs are absorbed by arbitrage capital; the market remains in a "no-run" region of the equilibrium correspondence. Once $\sigma_t$ exceeds $\sigma^*$, the sign of the marginal effect flips, and increases in adverse selection costs are associated with declines in open interest, consistent with a coordinated withdrawal of arbitrage capital. In our empirical work, we estimate (7) separately for distinct stress episodes (Terra, FTX, SVB, and the 2021 China ban phases) and compute $\sigma^*$. We then compare the location of $\sigma^*$ relative to realized volatility and the behavior of open interest across episodes.

## 4.3 Arbitrage Mechanism

We examine how closely the perpetual funding rate tracked the basis around the stress episodes. Our analysis centers on the fundamental arbitrage relationship that governs perpetual contract pricing, where funding payments serve as the primary mechanism for anchoring derivative prices to their underlying spot values. To understand the arbitrage mechanism, we examine the funding rates during stress regimes. The funding rate consists of two components: a fixed interest rate and a dynamic component (premium) that adjusts to the basis. The variation in the dynamic component captures the efficacy of arbitrage mechanism. Funding-rate

elasticity captures how closely the dynamic component of the funding rate adjusts to the basis, and is computed as the coefficient from regressing funding rates on the basis. Since the interest rate is fixed, the coefficient of regression captures the impact of variations in dynamic rates on the basis. A theoretical elasticity of one implies that the arbitrage mechanism fully aligns the perpetual price with the spot price and effectively closes the perp-spot gap.

$$f_t = f^{\text{fix}} + f_t^{\text{dyn}}. \tag{11}$$

$$\frac{\partial f_t}{\partial t} = \frac{\partial f_t^{\text{dyn}}}{\partial t}, \tag{12}$$

Following (He et al., 2022) to quantify the degree of futures–spot mispricing and assuming $r - r' \approx 0$ (a reasonable approximation since there is no interest on ETH holdings), the annualised no-arbitrage deviation measure in ETH, $\rho$, for each minute can be computed as

$$\rho_t \approx \kappa \left(\log F_t - \log S_t\right).$$

where $F_t$ denotes the perpetual futures price and $S_t$ represents the spot index value at time t. The scaling factor emerges from the institutional structure of cryptocurrency perpetual markets in Kraken, where funding payments typically occur at hourly intervals (three times daily). The annualization thus requires:

$$\kappa = \underbrace{24}_{} \text{ payments/day} \times \underbrace{365}_{} \text{ days/year} = 8760 \text{ hours/year} \tag{13}$$

This scaling transforms minute-level basis observations into annualized percentage rates, facilitating economic interpretation. We utilize minute-by-minute Ethereum/USD perpetual futures data from Kraken, spanning 2020 through

2023, encompassing multiple stress episodes. Kraken's hourly funding (AsiaNext Exchange, 2024) provides an ideal laboratory for examining funding-basis dynamics at high frequency. The dataset includes the perpetual futures price $F_t$, the corresponding spot index price $S_t$, the funding rate $f_t$ applied to positions, as well as trading activity metrics like open interest, bid–ask spreads, and volume. Following exchange conventions, the hourly funding rate $f_t$ is settled at the end of each hour, while the contemporaneous basis is defined as the log difference between mark and index prices. The mark price is the exchange's internal fair-value estimate of the perpetual contract. The index price represents the reference spot value of ETH/USD.

$$b_t = \ln(\text{mark}_t) - \ln(\text{index}_t). \tag{14}$$

We construct an evenly spaced one-hour panel and compute the average premium over the preceding hour,

$$\bar{b}_{t-1:t} = \frac{1}{60} \sum_{s \in [t-1,t)} b_s, \tag{15}$$

which serves as the explanatory variable for the next-hour funding payment. Both funding and basis are expressed in annualized units by multiplying by $\kappa = 24 \times 365 = 8760$. The baseline regression is

$$f_t^{APR} = \alpha + \beta \rho_t^{APR} + \varepsilon_t, \tag{16}$$

where $\rho_t^{APR} = \kappa \bar{b}_{t-1:t}$ represents the annualized premium, and $\beta$ captures the elasticity of the dynamic component of the funding rate, $\beta = \frac{\partial f_t^{\text{dyn}}}{\partial \rho_t^{APR}}$.

The coefficient $\beta$ captures the funding-basis elasticity—the percentage change in funding rates associated with a one percentage point change in the annualized basis. In a frictionless benchmark with fully effective arbitrage, we would expect

$\beta = 1$, meaning funding adjusts one-for-one with the basis and fully transmits price discrepancies into funding incentives. Values of $\beta$ below 1 indicate underreaction: funding does not adjust enough to close the basis, suggesting a weakened or constrained arbitrage channel. Values of $\beta$ above 1 indicate overreaction: funding moves more than one-for-one relative to the basis, consistent with an aggressive response in which the funding mechanism strongly rewards (or penalizes) positions and draws in additional arbitrage activity.

.

# 5 Results

## 5.1 Baseline estimates

We begin by examining the evolution of the dynamic component of funding basis elasticity ($\beta$) for Ethereum perpetual contracts on Kraken from **September 2020 to April 2022**. This interval includes multiple market cycles, volatility regimes, and systemic stress events. We examine the funding–basis relation during normal market conditions to establish a baseline against which stress-period distortions can be measured. Establishing this benchmark is essential: it provides clear evidence of what "efficient" arbitrage looks like during calm market periods.Figure 1b displays the scaled elasticity series across this 20-month window. The subsequent stress events—Terra's collapse, the FTX bankruptcy, and the SVB collapse can therefore be interpreted as deviations from this baseline. Table 2 presents baseline results for arbitrage functionality. We find that, in the calm periods, the mapping between the prior-hour premium and next-hour funding rate is stable and consistent, with funding payments from long to short most of the periods. This behavior is fully consistent with the functioning arbitrage mechanism during calm periods.

## 5.2 Funding Rate Dynamics during Market Stress

Do all market stress limit *"arbitrage"* in perpetual futures? We examine funding rates, a crucial *perp spot* price-convergence mechanism, during periods of extreme market stress. Figure 2 presents the 48-hour rolling average funding rate (annualized) across three distinct periods: the Pre-Crash, Crash, and Post-Crash periods, shaded in green, red, and blue, respectively. This is a period of stress transmission from the Terra, FTX, and SVB ecosystems to ETH derivatives markets. Funding rate remained stable pre-crash but shows a steep decline once the stress starts.

Next, we examine how closely the perpetual funding rate tracked basis around stress episodes. Understanding when and why funding elasticity breaks down can reveal deeper insights into the behavior of traders and the health of derivatives markets. Figure 3 shows the hourly relation between the next-hour funding rate and the prior-hour premium.

The baseline regression is.

$$f_t^{APR} = \alpha + \beta \rho_t^{APR} + \varepsilon_t, \tag{17}$$

where $\rho_t^{APR} = \kappa \bar{b}_{t-1:t}$ represents the annualized premium.

Figure (3, and 4) and Table 3 present results. Results suggest that the pre-Terra period exhibits $\beta = 1.44 \times 10^{-5}$ $(R^2 = 0.113)$, corresponding to a scaled elasticity of $0.64$. During the Terra crash (May 7–10), $\beta$ is insignificant and declines to $0.64 \times 10^{-5}$ $(R^2 = 0.11)$, with a scaled elasticity of $0.38$. An increase in the intercept [1.44 (pre) to 8.41(post)] indicates an increased response of the dynamic component of the funding rate. This suggests that arbitrageurs were attempting to close the basis; however, a decrease in elasticity ($\beta$) indicates that the arbitrage mechanism was ineffective in restoring the basis. In the post-crash phase (May 11–30), $\beta$ increases

to $1.52 \times 10^{-5}$ ($R^2 = 0.13$), and the scaled elasticity rises to $0.91$, consistent with restoration of arbitrage mechanism.

The pre-FTX period exhibits $\beta = 2.43 \times 10^{-5}$ ($R^2 = 0.64$), corresponding to a scaled elasticity of $1.44$, indicating a strong and stable link between funding rates and the basis. The intercept is negative ($\alpha = -6.95 \times 10^{-6}$), implying that, on average, perpetual prices trade below spot and shorts pay longs even after controlling for the basis. During the FTX crash window (November 6–11), $\beta$ increases to $2.91 \times 10^{-5}$ ($R^2 = 0.75$), with a scaled elasticity of $1.72$, reflecting a temporary amplification of the pass-through from basis deviations to funding adjustments. At the same time, the intercept becomes more negative ($\alpha = -1.39 \times 10^{-5}$), indicating that the level of funding shifts further downward and that short positions pay a larger amount to longs during the crash. In the post-crash phase (November 12–December 15), $\beta$ rises further to $3.18 \times 10^{-5}$ ($R^2 = 0.91$), and the scaled elasticity reaches $1.88$, consistent with a strengthened funding–basis linkage as markets stabilized, while the intercept remains negative ($\alpha = -9.97 \times 10^{-6}$), signaling a persistently discounted perp relative to spot and continued net payments from shorts to longs. Funding rate responsiveness (Figure **??**) suggests effective arbitrage and strong funding–basis correction.

Silicon Valley Bank faced a rapid deposit run on March 8- 9, 2023, and was closed by regulators on March 10, 2023. Withdrawal requests of around $42B (25% of deposits) hit within eight hours; end-of-day cash balance about –$958M. Funding $\alpha$ turned more negative as balance sheets tightened and shorts demanded premia, but $\beta$ did not break, so the arbitrage mapping still functioned—like FTX. Around this failure, ETH perpetuals on Kraken exhibit a *tight funding–basis relation with a downward level shift*: the crash-window scatter shows more negative funding for a given premium, while the elasticity remains near or above its pre-period level and only eases afterward. This mirrors the FTX episode in which arbitrage effectiveness is preserved. By contrast, during Terra's collapse, arbitrage effectiveness collapses.

## 5.3 Capital Constraints and Global Games Approach

Why would similarly devastating market events have opposite effects on funding rate elasticity? To understand the source of this asymmetry, we view open interest as the stock of arbitrage capital deployed in perpetuals to correct the deviations in basis and interpret its collapse as a run on arbitrage capital. We find that open interest behaves very differently across crash episodes. During the Terra collapse, average ETH perpetual open interest falls from about $58.7 million in the pre period to $20.3 million in the post period—a decline of roughly 65%—with the crash window itself exhibiting both elevated volatility and extremely wide dispersion in open interest. By contrast, during the FTX bankruptcy and the SVB failure, ETH perpetual average open interest remains in a relatively narrow range: it drifts slightly *upward* from $10.9m to $11.8m across the FTX episode and declines by only about 1% (from $9.38m to $9.28m) across the SVB window. In other words, only during the Terra crash period, there was a run on the ETH-perpetual market as seen by capital constraints to arbitrage.

To understand the cause of the run only during the Terra crash episode, we bring to this setting a global-games perspective in the spirit of Morris and Shin (2001); Goldstein and Pauzner (2005). In global-games models, agents decide whether to stay or run based on a fundamental state and noisy private signals about that state, and a threshold emerges at which the unique equilibrium switches from no-run to run. We adapt this framework (Equation 10) by modeling arbitrageurs' stay-or-exit decisions (change in open interest) as a function of public stress (volatility) and private signals (adverse selection costs). Results from Table 7 reveal a sharp contrast. During the Terra crash, we estimate $\beta_{AS} = 4.591$ (p<0.01) and $\beta_{AS\times\sigma} = -257.249$ (p<0.01), yielding a volatility threshold $\sigma^* = 1.785\%$ per hour. This threshold lies within the range of realized hourly volatility during the crisis (Fig 5, Terra crash median $\sigma = 0.71\%$, with peaks exceeding $1.8\%$ for 37 hours). As $\sigma_t$ crossed $\sigma^*$,

the marginal effect of adverse selection on open interest flipped from positive to negative, triggering a coordinated withdrawal of arbitrage capital—a classic global-games run equilibrium. Open interest collapsed by 65% (Figure 8a), consistent with a capital-destructive regime. During the pre and post-crash periods, volatility remained below the threshold all the time, and open interest remained stable.

In stark contrast, during the FTX and SVB crash episodes, the global-games mechanism remained dormant. For FTX, $\beta_{AS\times\sigma}$ is statistically insignificant ($-6.062$, p>0.10), implying no meaningful threshold. For SVB, the interaction term is positive during the crisis ($85.659$, p>0.10), but the opposite sign is required for a coordination threshold. Consequently, even though volatility crossed the threshold in both events (FTX crash $\sigma = 0.9\%$, SVB crash $\sigma = 1.1\%$), adverse selection costs did not induce the same run dynamics. Moreover, during the SVB Collapse period, the threshold volatility of 1.1% was statistically significant post-crash, and volatility remained below this threshold before. Open interest remained stable (FTX: +36%; SVB: -4%), and the funding–basis elasticity stayed near or above 1 (Table 4), indicating preserved arbitrage functionality.

# 6 Robustness

The phased China 2021 regulatory crackdown is a powerful natural experiment to test the limits of arbitrage under the global games framework. The ban is correctly seen, as recent work shows, as one of the sharpest regulatory shocks in crypto's history, causing a market-wide collapse in prices, a spike in volatility, and—most relevant for this study—plummeting market quality through mass liquidity withdrawal and forced exit of arbitrageurs. (Manwaring, 2021).

China's regulators progressively escalated restrictions on crypto activity from May to September 2021. On May 18, financial associations ordered banks and

payment firms to halt all crypto-related services, citing financial-stability risks. On May 21, the State Council explicitly targeted Bitcoin mining and trading, triggering nationwide enforcement and shutdown orders across major mining provinces through June. By September 24, the PBoC and nine national agencies declared all crypto transactions illegal and barred foreign exchanges from serving Chinese users. The phased nature of regulatory announcements created an information-revelation sequence: initial policy ambiguity (Phase I) was gradually resolved through enforcement actions (Phases II–III) and culminating in a comprehensive, symmetric ban (Phase IV) as shown in Table 8.

We first show that the funding rate elasticity collapses towards zero (Fig 10 and Table 9 ), in Phase I, with the highest level of information asymmetry. During the pre-crackdown period, the funding mechanism operated with near-ideal efficiency ($\beta_{\text{scaled}} = 1.48$), where 30% of funding-rate variation was explained by observable basis deviations, consistent with well-functioning arbitrage markets. In Phase I, both funding elasticity collapses to 0; $\beta_{\text{scaled}} = 0.09$, $R^2 \approx 0$). The estimates indicate that the dynamic component of the funding rate ($f_t^{\text{dyn}}$) stops responding to mispricing. Elasticity recovers towards a value of one post Phase I period.

To understand the collapse in funding elasticity around the China crackdown, we examine open interest as it represents arbitrage capital. A sudden decline of open interest implies a run on perpetual. (Figure 12). Prior to the ban, average open interest is about \$72.7 million, with values ranging from \$45.7 million to more than \$102 million. During Phase I of the ban, open interest falls to \$39.4 million—a decline of roughly 45% relative to the pre-ban mean—and the lower tail reaches levels near \$20 million. In Phases II and III, open interest stabilizes at an even lower plateau of \$27–31 million, before partially recovering in Phase IV (mean \$77.3 million). Thus, the China ban episode features a discrete, policy–driven run on ETH perpetual positions in Phase I, followed by a prolonged regime of structurally lower arbitrage capital, rather than a short–lived liquidity shock.

We adapt the global games approach (Equation 10) to understand the cause of the collapse of open interest during Phase I. We model arbitrageurs' stay-or-exit decisions (change in open interest) as a function of public stress (volatility) and private signals (adverse selection costs). Results are presented in Table 11 and Figure 11. During Phase I, we estimate $\beta_{AS} = 1.987$ (p<0.01) and $\beta_{AS\times\sigma} = -42.837$ (p<0.01), yielding a volatility threshold $\sigma^* = 4.6\%$ per hour. This threshold is statistically significant and lies within the range of realized hourly volatility during Phase I (Fig. 11, Terra crash median $\sigma = 0.71\%$, with peaks exceeding $1.8\%$ for 37 hours). As $\sigma_t$ crossed $\sigma^*$, the marginal effect of adverse selection on open interest flipped from positive to negative, triggering a coordinated withdrawal of arbitrage capital—a classic global-games run equilibrium. Open interest collapsed by 45% (Figure 12) soon after volatility crossed the threshold $\sigma^*$ during Phase I. During other Phases, $\sigma^*$ was statistically insignificant, and the run was not triggered despite volatility $\sigma$ crossing the threshold $\sigma^*$ during Phases III and IV. Results from Phase I (Table 9 and Figure 10) suggest that the funding rate could not correct the basis, and arbitrage broke down. Overall, these robustness checks confirm that the results are fully consistent with the patterns observed during the Terra, FTX, and SVB episodes, and they extend the global-games framework of (Goldstein and Pauzner, 2005; Morris and Shin, 2001) and the limits-to-arbitrage insights of (Shleifer and Vishny, 1997) to an empirical setting in perpetual futures markets.

# 7 Conclusion

This paper applies the global games framework to identify when and why arbitrage mechanisms fail in derivatives markets. Our contribution to the literature is threefold. First, we document a new empirical puzzle: crises of comparable severity—Terra ($60B), FTX ($32B), and SVB ($42B)—produce starkly opposite

effects on arbitrage functionality. This finding suggests that crisis magnitude alone may not fully determine arbitrage capacity, pointing to additional mechanisms at work. Second, we provide a theoretical resolution grounded in the global games run threshold equilibrium model. By mapping the global games framework of Morris and Shin (2001) and Goldstein and Pauzner (2005) onto perpetual futures markets—with spot volatility as a public signal, adverse selection costs as private signals, and open interest as the aggregate coordination outcome—we derive an endogenous volatility threshold that separates regimes from capital-constrained regimes – run equilibria. This threshold is empirically estimable and statistically significant only during Terra's collapse, precisely when arbitrage failed.

Third, we bridge two foundational literatures. We extend classical capital-based limits to arbitrage by showing that constraints can emerge endogenously through coordination failure, complementing capital shortage as an explanatory channel. Simultaneously, we provide an empirical application of global games coordination thresholds to the derivatives (perpetual futures) market, by identifying the specific mechanism through which arbitrage breaks down under stress. Our findings carry implications for exchange design, regulatory surveillance, and risk management: monitoring volatility relative to estimated thresholds may provide early warning of arbitrage mechanism failure before capital flight becomes visible.

# References


Ackerer, D., Hugonnier, J., and Jermann, U. (2024). Perpetual futures pricing. Technical report, National Bureau of Economic Research.

Alexander, C., Choi, J., Park, H., and Sohn, S. (2020). Bitmex bitcoin derivatives: Price discovery, informational efficiency, and hedging effectiveness. *Journal of Futures Markets*, 40(1):23–43.

Angeris, G., Chitra, T., Evans, A., and Lorig, M. (2023). A primer on perpetuals. *SIAM Journal on Financial Mathematics*, 14(1):SC17–SC30.

AsiaNext Exchange (2024). Asianext crypto futures methodology overview. https://www.asianext.com/wp-content/uploads/2025/01/AsiaNext-Crypto-Futures-Methodology-Overview-Dec-2024.pdf. Accessed: YYYY-MM-DD.

Brunnermeier, M. K. and Pedersen, L. H. (2005). Predatory trading. *The Journal of Finance*, 60(4):1825–1863.

De Blasis, R. and Webb, A. (2022). Arbitrage, contract design, and market structure in bitcoin futures markets. *Journal of Futures Markets*, 42(3):492–524.

Du, W., Tepper, A., and Verdelhan, A. (2018). Deviations from covered interest rate parity. *The Journal of Finance*, 73(3):915–957.

Glosten, L. R. and Harris, L. E. (1988). Estimating the components of the bid/ask spread. *Journal of financial Economics*, 21(1):123–142.

Goldstein, I. and Pauzner, A. (2005). Demand–deposit contracts and the probability of bank runs. *the Journal of Finance*, 60(3):1293–1327.

Gornall, W., Rinaldi, M., and Xiao, Y. (2024). Funding payments crisis-proofed bitcoin's perpetual futures. *Available at SSRN*.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408.

Hasbrouck, J. (1993). Assessing the quality of a security market: A new approach to transaction-cost measurement. *The Review of Financial Studies*, 6(1):191–212.

He, S., Manela, A., Ross, O., and von Wachter, V. (2022). Fundamentals of perpetual futures. *arXiv preprint arXiv:2212.06888*.

Joshi, R. (2025). Essays on stablecoins'stability.

Kim, J. and Park, H. (2025). Designing funding rates for perpetual futures in cryptocurrency markets. *arXiv preprint arXiv:2506.08573*.

Liu, J., Makarov, I., and Schoar, A. (2023). Anatomy of a run: The terra luna crash. Technical report, National Bureau of Economic Research.

Lyons, R. K. and Viswanath-Natraj, G. (2023). What keeps stablecoins stable? *Journal of International Money and Finance*, 131:102777.

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

Manwaring, K. (2021). The volatility implications of the chinese cryptocurrency ban. Graduate student research report, Utah State University. Accessed: Month Day, 2025.

Morris, S. and Shin, H. S. (2001). Global games: Theory and applications.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of finance*, 39(4):1127–1139.

Shiller, R. J. (1993). Measuring asset values for cash settlement in derivative markets: hedonic repeated measures indices and perpetual futures. *The Journal of Finance*, 48(3):911–931.

Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *The Journal of finance*, 52(1):35–55.

Streltsov, A. and Ruan, Q. (2022). Perpetual price discovery and crypto market quality. In *Perpetual Price Discovery and Crypto Market Quality: Streltsov, Artem | uRuan, Qihong*. [Sl]: SSRN.

**Table 1:** Summary Statistics — Ethereum

**Panel A: Terra**

| Variable | Pre Mean | Pre SD | Crash Mean | Crash SD | Post Mean | Post SD |
|---|---|---|---|---|---|---|
| Price | 2913.47 | 120.30 | 2256.09 | 216.47 | 1951.43 | 100.50 |
| Basis | 0.17 | 0.43 | 0.32 | 0.54 | -0.05 | 0.25 |
| Funding Rate | 0.002 | 0.005 | 0.01 | 0.01 | -0.002 | 0.01 |
| Spread | 1.72 | 0.37 | 1.69 | 0.61 | 1.18 | 0.21 |
| AS_cost_10k | 0.27 | 1.84 | 1.84 | 2.45 | 0.36 | 2.05 |
| Open Interest | 58.74 | 1.14 | 49.27 | 18.68 | 20.30 | 1.12 |
| $\sigma_{ETH}$ | 0.01 | 0.002 | 0.02 | 0.01 | 0.01 | 0.002 |
| Order Flow | -0.002 | 0.77 | -0.07 | 0.61 | 0.08 | 0.74 |
| Momentum | -0.0003 | 0.01 | -0.002 | 0.02 | -0.0001 | 0.01 |

**Panel B: FTX**

| Variable | Pre Mean | Pre SD | Crash Mean | Crash SD | Post Mean | Post SD |
|---|---|---|---|---|---|---|
| Price | 1442.30 | 136.87 | 1376.60 | 167.36 | 1234.27 | 50.08 |
| Basis | -0.11 | 0.28 | -0.85 | 0.89 | -0.47 | 0.71 |
| Funding Rate | -0.003 | 0.01 | -0.02 | 0.02 | -0.01 | 0.02 |
| Spread | 0.63 | 0.25 | 1.30 | 0.60 | 0.73 | 0.22 |
| AS_cost_10k | 0.68 | 2.35 | 1.40 | 1.50 | 0.18 | 0.58 |
| Open Interest | 10.96 | 0.29 | 11.33 | 1.11 | 11.83 | 1.19 |
| $\sigma_{ETH}$ | 0.01 | 0.003 | 0.02 | 0.01 | 0.005 | 0.003 |
| Order Flow | 0.06 | 0.76 | -0.04 | 0.62 | -0.01 | 0.70 |
| Momentum | 0.0004 | 0.01 | -0.002 | 0.02 | 0.0000 | 0.01 |

**Panel C: SVB**

| Variable | Pre Mean | Pre SD | Crash Mean | Crash SD | Post Mean | Post SD |
|---|---|---|---|---|---|---|
| Price | 1630.40 | 48.99 | 1484.58 | 61.26 | 1534.51 | 83.52 |
| Basis | 0.16 | 0.42 | -0.06 | 0.69 | -2.65 | 1.94 |
| Funding Rate | 0.002 | 0.01 | -0.0003 | 0.01 | -0.05 | 0.04 |
| Spread | 0.43 | 0.13 | 0.54 | 0.19 | 0.81 | 0.20 |
| AS_cost_10k | 0.07 | 0.27 | 0.18 | 0.32 | 0.22 | 0.39 |
| Open Interest | 9.36 | 0.28 | 9.05 | 0.29 | 8.94 | 0.44 |
| $\sigma_{ETH}$ | 0.005 | 0.003 | 0.01 | 0.004 | 0.01 | 0.002 |
| Order Flow | 0.10 | 0.74 | -0.25 | 0.63 | 0.04 | 0.55 |
| Momentum | 0.0000 | 0.01 | 0.001 | 0.01 | 0.001 | 0.01 |

**Table 2: Funding–Basis Elasticity in Normal Times (Nov 2021 – Apr 2022).** This table reports monthly regressions of the hourly funding rate ($y_t$) on the prior-hour basis premium ($x_t$): $y_t = \alpha + \beta x_t + \varepsilon_t$. Funding and basis are annualized (APR). Coefficients $\alpha$ and $\beta$ are reported in units of $10^{-5}$. Scaled elasticity $\beta_{\text{scaled}}$ normalizes $\beta$ relative to full pass-through.

| | **Nov 2021** | **Dec 2021** | **Jan 2022** | **Feb 2022** | **Mar 2022** | **Apr 2022** |
|---|---|---|---|---|---|---|
| $\alpha$ ($\times 10^{-5}$) | $1.30^{***}$ | $1.04^{***}$ | $-0.501^{**}$ | $1.29^{***}$ | $-0.445^{**}$ | $1.56^{***}$ |
| | $(6.55)$ | $(7.31)$ | $(-2.21)$ | $(6.03)$ | $(-2.19)$ | $(9.60)$ |
| $\beta$ ($\times 10^{-5}$) | $1.34^{***}$ | $0.753^{***}$ | $1.32^{***}$ | $0.990^{***}$ | $1.14^{***}$ | $0.856^{***}$ |
| | $(17.29)$ | $(9.35)$ | $(10.68)$ | $(7.93)$ | $(9.48)$ | $(7.87)$ |
| $\beta_{\text{scaled}}$ | 0.79 | 0.45 | 0.78 | 0.59 | 0.68 | 0.51 |
| $R^2$ | 0.301 | 0.108 | 0.137 | 0.088 | 0.111 | 0.082 |
| $N$ | 697 | 721 | 721 | 649 | 721 | 697 |

*Note:* $\beta_{\text{scaled}} = \beta / \text{median}(f_t/\rho_t)$ within each window. Stars denote significance at 10% (*), 5% (**), and 1% (***). T-statistics are reported in parentheses.

**Table 3: Funding–Basis Elasticity Across Crisis Episodes.** This table reports hourly regressions of next-hour funding ($y_t$) on the prior-hour basis premium ($x_t$): $y_t = \alpha + \beta x_t + \varepsilon_t$. Funding and basis are annualized (APR). Coefficients $\alpha$ and $\beta$ are expressed in units of $10^{-5}$. Scaled elasticity $\beta_{\text{scaled}} \equiv \beta/\text{median}(f_t/\rho_t)$ normalizes $\beta$ so that $\beta_{\text{scaled}} \approx 1$ corresponds to full pass-through.

| | **Terra Collapse** | | | **FTX Collapse** | | | **SVB Episode** | | |
|---|---|---|---|---|---|---|---|---|---|
| | Pre | Crash | Post | Pre | Crash | Post | Pre | Crash | Post |
| $\alpha$ ($\times 10^{-5}$) | $1.44^{***}$ | $8.41^{***}$ | $-0.64^{**}$ | $-0.695^{***}$ | $-1.39$ | $-0.997^{***}$ | $0.235$ | $1.56^{***}$ | $-2.82^{**}$ |
| | $(7.98)$ | $(9.31)$ | $(-2.26)$ | $(-4.28)$ | $(-1.05)$ | $(-4.44)$ | $(1.54)$ | $(2.76)$ | $(-2.01)$ |
| $\beta$ ($\times 10^{-5}$) | $1.08^{***}$ | $0.643$ | $1.53^{***}$ | $2.43^{***}$ | $2.91^{***}$ | $3.18^{***}$ | $2.13^{***}$ | $1.71^{***}$ | $2.33^{***}$ |
| | $(8.18)$ | $(1.62)$ | $(8.23)$ | $(29.93)$ | $(18.79)$ | $(88.87)$ | $(33.34)$ | $(5.66)$ | $(20.90)$ |
| $\beta_{\text{scaled}}$ | 0.64 | 0.38 | 0.91 | 1.44 | 1.72 | 1.88 | 1.26 | 1.01 | 1.38 |
| $R^2$ | 0.113 | 0.036 | 0.130 | 0.641 | 0.748 | 0.907 | 0.689 | 0.406 | 0.725 |
| $N$ | 528 | 73 | 456 | 504 | 121 | 816 | 504 | 49 | 168 |

*Notes:* T-statistics in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. $\beta_{\text{scaled}}$ uses a common empirical scale factor $\text{median}(f_t/\rho_t) = 1.69 \times 10^{-5}$ for comparability across windows. *Event windows:* Pre Terra (Apr 15–May 6), Terra Crash (May 7–10), Post Terra (May 11–30); Pre FTX (Oct 15–Nov 5), Crash FTX (Nov 6–11), Post FTX (Nov 11–Dec 15); Pre SVB (Feb 15–Mar 7), SVB Crash (Mar 8–10), Post SVB (Mar 11–18).

**Table 4: Adverse Selection Summary Statistics by Event**

| Statistic | Terra | FTX | SVB |
|---|---|---|---|
| *PRE Median* | 0.185 | 0.171 | 0.111 |
| *PRE SD* | 3.495 | 1.242 | 0.558 |
| *PRE P25* | -0.000 | -0.029 | -0.116 |
| *PRE P75* | 0.392 | 0.416 | 0.395 |
| *N* | 455 | 941 | 363 |
| *CRASH Median* | 0.313 | 0.329 | 0.161 |
| *CRASH SD* | 0.339 | 0.564 | 0.503 |
| *CRASH P25* | 0.122 | 0.168 | -0.082 |
| *CRASH P75* | 0.586 | 0.586 | 0.357 |
| *N* | 123 | 73 | 39 |
| *POST Median* | 0.192 | 0.140 | 0.162 |
| *POST SD* | 0.692 | 0.681 | 0.315 |
| *POST P25* | -0.050 | -0.046 | 0.041 |
| *POST P75* | 0.458 | 0.372 | 0.331 |
| *N* | 364 | 824 | 96 |

**Table 5: Adverse Selection Cost Summary Statistics for Ethereum**
Summary statistics are reported for the pre-crisis, crash, and post-crisis windows of the Terra, FTX, and SVB episodes. Adverse selection costs, realized in next hour $t$, are computed using the median adverse selection estimate between consecutive hours and the corresponding 24-hour trading volume. Specifically, for each hour $t$, we compute

$$\text{AS Cost}_t = \text{RawAS}_{[t-1 \to t]} \times \left( \frac{\text{Volume}_t^{24h}}{\text{Price}_t} \right),$$

where $\text{RawAS}_{[t-1 \to t]}$ denotes the median adverse-selection component between $t-1$ and $t$ (in USD/ETH), $\text{Volume}_t^{24h}$ is the 24-hour notional trading volume (in USD), and $\text{Price}_t$ is the spot ETH price (USD/ETH). This yields the total dollar value of informational trading costs in each hourly window.

| **Statistic** | **Terra** | **FTX** | **SVB** |
|---|---|---|---|
| *PRE Median* | 1433.96 | 1988.55 | 485.81 |
| *PRE SD* | 18638.99 | 19967.02 | 2721.63 |
| *PRE P25* | -6.20 | -277.39 | -389.21 |
| *PRE P75* | 3543.42 | 7196.47 | 1718.77 |
| *PRE N* | 455 | 919 | 363 |
| *CRASH Median* | 12911.05 | 14561.84 | 455.97 |
| *CRASH SD* | 24691.19 | 17277.00 | 1319.94 |
| *CRASH P25* | 2775.21 | 5512.76 | -228.98 |
| *CRASH P75* | 27179.43 | 19992.39 | 990.32 |
| *CRASH N* | 123 | 73 | 39 |
| *POST Median* | 3327.19 | 979.86 | 1826.78 |
| *POST SD* | 20450.48 | 5685.78 | 3894.06 |
| *POST P25* | -848.34 | -290.99 | 504.67 |
| *POST P75* | 9769.68 | 3018.78 | 3344.97 |
| *POST N* | 363 | 824 | 96 |

**Table 6:** Open Interest Response to Market Forces

| | Pre-Crisis | Crisis | Post-Crisis |
|---|---|---|---|
| **Panel A: Terra Collapse (May 2022)** | | | |
| AS_cost$_{t-1}$ | $-0.040$ | $4.591^{***}$ | $0.174$ |
| | $(0.141)$ | $(1.286)$ | $(0.132)$ |
| Realized Volatility ($\sigma_{ETH,t-1}$) | $-133.144^{***}$ | $-370.142^{**}$ | $-11.130$ |
| | $(24.292)$ | $(152.409)$ | $(26.116)$ |
| Momentum ($Mom_{t-1}$) | $-5.861$ | $-66.754$ | $-4.186$ |
| | $(8.886)$ | $(45.946)$ | $(7.056)$ |
| Order Flow Imbalance ($OF_{i,t-1}$) | $-0.085$ | $0.274$ | $0.171^{*}$ |
| | $(0.076)$ | $(1.617)$ | $(0.092)$ |
| AS_cost$_{t-1} \times \sigma_{ETH,t-1}$ | $27.073$ | $-257.249^{***}$ | $-27.683^{*}$ |
| | $(33.698)$ | $(50.142)$ | $(14.760)$ |
| Observations | 462 | 122 | 387 |
| Hourly FE | Yes | Yes | Yes |
| $R^2$ | 0.088 | 0.640 | 0.073 |
| **Panel B: FTX Collapse (Nov 2022)** | | | |
| AS_cost$_{t-1}$ | $-0.026^{*}$ | $0.052$ | $-0.152^{*}$ |
| | $(0.015)$ | $(0.175)$ | $(0.085)$ |
| Realized Volatility ($\sigma_{ETH,t-1}$) | $-3.515$ | $71.399^{***}$ | $-24.398^{**}$ |
| | $(5.184)$ | $(21.686)$ | $(10.569)$ |
| Momentum ($Mom_{t-1}$) | $0.782$ | $-1.053$ | $6.114$ |
| | $(2.189)$ | $(8.149)$ | $(4.266)$ |
| Order Flow Imbalance ($OF_{i,t-1}$) | $0.020$ | $-0.119$ | $0.078$ |
| | $(0.022)$ | $(0.272)$ | $(0.064)$ |
| AS_cost$_{t-1} \times \sigma_{ETH,t-1}$ | $1.726$ | $-6.062$ | $6.698$ |
| | $(1.728)$ | $(8.835)$ | $(6.516)$ |
| Observations | 316 | 73 | 446 |
| $R^2$ | 0.079 | 0.431 | 0.049 |
| Hourly FE | Yes | Yes | Yes |

*Table continued on next page.*

Table 6 (continued)

| | Pre-Crisis | Crisis | Post-Crisis |
|---|---|---|---|
| **Panel C: SVB Collapse (Mar 2023)** | | | |
| AS_cost$_{t-1}$ | 0.171 | 0.122 | 1.751*** |
| | (0.146) | (0.393) | (0.541) |
| Realized Volatility ($\sigma_{ETH,t-1}$) | 9.743 | −78.825*** | 25.132 |
| | (6.249) | (16.025) | (28.990) |
| Momentum ($Mom_{t-1}$) | 4.564* | −3.945 | 11.393*** |
| | (2.561) | (2.230) | (3.801) |
| Order Flow Imbalance ($OF_{i,t-1}$) | −0.023 | −0.034 | −0.018 |
| | (0.023) | (0.028) | (0.091) |
| AS_cost$_{t-1} \times \sigma_{ETH,t-1}$ | −27.769 | 85.659 | −157.659*** |
| | (21.123) | (98.908) | (54.488) |
| Observations | 345 | 39 | 96 |
| $R^2$ | 0.064 | 0.948 | 0.285 |
| Hourly FE | Yes | Yes | Yes |

Notes: This table reports the impact of market frictions on Ethereum perpetuals open interest during different stress episodes. The dependent variable is hourly open interest (in millions USD). This table reports volatility thresholds estimated from the regression: $OI_t = \alpha + \beta_{AS} \cdot \text{AS\_cost}_{t-1} + \beta_\sigma \cdot \sigma_{t-1} + \beta_{AS\times\sigma} \cdot \text{AS\_cost}_{t-1} \times \sigma_{t-1} + \gamma' X_{t-1} + \text{Hour FE} + \varepsilon_t$. Pre-crisis periods span 30 days before each event, crisis windows are defined as: Terra (May 8-12, 2022), FTX (November 6-9, 2022), SVB (March 8-9, 2023). Pre and Post-crisis periods extend 30 days before and after each event.

**Table 7: Volatility Threshold Estimates Across Crash Episodes**

This table reports volatility thresholds estimated from the regression: $OI_t = \alpha + \beta_{AS} \cdot \text{AS\_cost}_{t-1} + \beta_\sigma \cdot \sigma_{t-1} + \beta_{AS\times\sigma} \cdot \text{AS\_cost}_{t-1} \times \sigma_{t-1} + \gamma' X_{t-1} + \text{Hour FE} + \varepsilon_t$. The threshold volatility $\sigma^* = -\beta_{AS}/\beta_{AS\times\sigma}$ represents the level at which the marginal effect of adverse selection costs on open interest is zero. Run is determined based on whether the threshold was breached during any period and whether the coefficient pattern ($\beta_{AS} > 0$, $\beta_{AS\times\sigma} < 0$) indicates a coordination equilibrium. Only the Terra crash exhibits a statistically significant threshold (1.785% per hour) with the correct sign pattern and threshold breach, consistent with a global-games run equilibrium.

| **Crisis** | **Period** | $\beta_{AS}$ | $\beta_{AS\times\sigma}$ | **Threshold $\sigma^*$ (%/hour)** | **Run?** |
|---|---|---|---|---|---|
| | | | **Terra Collapse (May 2022)** | | |
| Terra | Pre | $-0.040$ | $27.073$ | 0.148 | No |
| | | (0.141) | (33.698) | | |
| Terra | **Crash** | $\mathbf{4.591}^{***}$ | $\mathbf{-257.249}^{***}$ | **1.785** | **Yes** |
| | | (1.286) | (50.142) | | |
| Terra | Post | $0.174$ | $-27.683^{*}$ | 0.628 | No |
| | | (0.132) | (14.760) | | |
| | | | **FTX Collapse (November 2022)** | | |
| FTX | Pre | $-0.026^{*}$ | $1.726$ | 1.5 | No |
| | | (0.015) | (1.728) | | |
| FTX | Crash | $0.052$ | $-6.062$ | 0.9 | No |
| | | (0.175) | (8.835) | | |
| FTX | Post | $-0.152^{*}$ | $6.698$ | 2.3 | No |
| | | (0.085) | (6.516) | | |
| | | | **SVB Collapse (March 2023)** | | |
| SVB | Pre | $0.171$ | $-27.769$ | 0.6 | No |
| | | (0.146) | (21.123) | | |
| SVB | Crash | $0.122$ | $85.659$ | -0.1 | No |
| | | (0.393) | (98.908) | | |
| SVB | Post | $\mathbf{1.751}^{***}$ | $\mathbf{-157.659}^{***}$ | 1.1 | **Yes** |
| | | (0.541) | (54.488) | | |

**Table 8:** China Ban Phases, Windows, and Adverse Selection (Median AS)

| Period | Window / Phase | Info Asymmetry | Median AS |
|---|---|---|---|
| 2021-03-22 – 2021-05-20 | 60-Day Pre | Low | 0.339 |
| 2021-05-21 – 2021-06-10 | Phase I | High | 0.712 |
| 2021-06-11 – 2021-06-25 | Phase II | Medium | 0.470 |
| 2021-06-26 – 2021-07-20 | Phase III | Medium | 0.416 |
| 2021-07-21 – 2021-09-24 | Phase IV | Declining | 0.330 |
| 2021-09-25 – 2021-11-23 | 60-Day Post | Low | 0.510 |

**Table 9: Funding–Basis Elasticity During China's Cryptocurrency Crackdown.** This table reports hourly regressions of the next-hour funding rate ($y_t$) on the prior-hour basis premium ($x_t$), $y_t = \alpha + \beta x_t + \varepsilon_t$. Funding and basis are annualized in APR. Coefficients $\alpha$ and $\beta$ are reported in units of $10^{-5}$. The scaled elasticity $\beta_{\text{scaled}}$ normalizes $\beta$ relative to $\beta = 1$.

| | **Pre** (Mar 22–May 20) | **Phase I** (May 21–Jun 10) | **Phase II** (Jun 11–Jun 25) | **Phase III** (Jun 26–Jul 20) | **Phase IV** (Jul 21–Sep 24) | **Post** (Sep 25–Nov 23) |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\alpha$ ($\times 10^{-5}$) | $14.1^{***}$ | $-0.347$ | $-3.91^{***}$ | $0.815^{**}$ | $4.03^{***}$ | $2.35^{***}$ |
| | $(16.62)$ | $(-0.46)$ | $(-5.91)$ | $(2.36)$ | $(17.78)$ | $(14.97)$ |
| $\beta$ ($\times 10^{-5}$) | $2.51^{***}$ | $0.156$ | $1.50^{***}$ | $0.896^{***}$ | $1.39^{***}$ | $1.30^{***}$ |
| | $(24.39)$ | $(0.94)$ | $(6.57)$ | $(5.13)$ | $(16.66)$ | $(21.34)$ |
| $\beta_{\text{scaled}}$ | 1.48 | 0.09 | 0.89 | 0.53 | 0.82 | 0.77 |
| $R^2$ | 0.296 | 0.002 | 0.114 | 0.044 | 0.151 | 0.244 |
| $N$ | 1417 | 458 | 337 | 577 | 1561 | 1417 |

*Note:* $\beta_{\text{scaled}}$ is computed as $\beta / \text{median}(f_t/\rho_t)$ within each window. *Event windows:* Pre (Mar 22–May 20), Phase I (May 21–Jun 10), Phase II (Jun 11–Jun 25), Phase III (Jun 26–Jul 20), Phase IV (Jul 21–Sep 24), Post (Sep 25–Nov 23).

**Table 10:** Open Interest Response to Adverse Selection Costs and Market Frictions Across China Ban Phases

This table reports the impact of market frictions on Ethereum perpetuals open interest during different stress episodes estimated from the regression: $OI_t = \alpha + \beta_{AS} \cdot \text{AS\_cost}_{t-1} + \beta_\sigma \cdot \sigma_{t-1} + \beta_{AS\times\sigma} \cdot \text{AS\_cost}_{t-1} \times \sigma_{t-1} + \gamma' X_{t-1} + \text{Hour FE} + \varepsilon_t$. The dependent variable is hourly open interest (in millions USD). Adverse selections costs are per 10k USD. Pre and Post-crisis periods extend 30 days before and after each event.

| | **Pre** | **PH1** | **PH2** | **PH3** | **PH4** | **Post** |
|---|---|---|---|---|---|---|
| AS_Cost_t-1 | 0.309** | 1.987*** | 0.098 | –0.054 | –0.185 | –0.211 |
| | (0.150) | (0.237) | (0.094) | (0.038) | (0.115) | (0.198) |
| Sigma_ETH_t-1 | 344.968*** | 609.467*** | 147.568*** | –52.156*** | –517.055*** | –1,076.180*** |
| | (103.238) | (144.538) | (35.370) | (19.432) | (66.920) | (95.455) |
| Mom_t-1 | 36.561 | –50.296 | –5.206 | –1.780 | –44.021** | –20.818 |
| | (35.971) | (46.107) | (12.050) | (6.015) | (22.244) | (36.145) |
| OF_i_t-1 | 0.900 | 1.538 | –0.441* | 0.006 | –0.404 | 0.415 |
| | (0.685) | (1.675) | (0.227) | (0.095) | (0.363) | (0.471) |
| AS_Cost_t-1 × Sigma_t-1 | 13.979 | –42.837*** | –2.502 | 6.198* | 26.362** | 39.400* |
| | (10.669) | (9.181) | (6.097) | (3.709) | (11.159) | (20.600) |
| Observations | 1,194 | 616 | 348 | 480 | 1,457 | 1,353 |
| $R^2$ | 0.079 | 0.164 | 0.129 | 0.038 | 0.050 | 0.092 |
| Hour FE | Y | Y | Y | Y | Y | Y |

**Table 11: Volatility Threshold Estimates — China Cryptocurrency Ban (2021)**

This table reports volatility thresholds estimated from the regression: $OI_t = \alpha + \beta_{AS} \cdot \text{AS\_cost}_{t-1} + \beta_\sigma \cdot \sigma_{t-1} + \beta_{AS\times\sigma} \cdot \text{AS\_cost}_{t-1} \times \sigma_{t-1} + \gamma' X_{t-1} + \text{Hour FE} + \varepsilon_t$. The threshold volatility $\sigma^* = -\beta_{AS}/\beta_{AS\times\sigma}$ represents the level at which the marginal effect of adverse selection costs on open interest is zero. Threshold volatility is defined as $\sigma^* = -\beta_{AS}/\beta_{AS\times\sigma}$. A "Run" is recorded only when: (i) $\beta_{AS} > 0$ and $\beta_{AS\times\sigma} < 0$ with $p < 0.01$, and (ii) the realized volatility in the event window crosses $\sigma^*$. Only Phase I satisfies both conditions, indicating a coordinated withdrawal of arbitrage capital during the China crypto ban regulatory period.

| **China Phase** | $\beta_{AS}$ | $\beta_{AS\times\sigma}$ | **Threshold $\sigma^*$ (%/hour)** | **Run?** |
|---|---|---|---|---|
| **PRE** | $0.309^{**}$ | $13.979$ | – | No |
| | (0.150) | (10.669) | | |
| **I** | $\mathbf{1.987}^{***}$ | $\mathbf{-42.837}^{***}$ | **4.6** | **Yes** |
| | (0.237) | (9.181) | | |
| **II** | $0.098$ | $-2.502$ | 3.9 | No |
| | (0.094) | (6.097) | | |
| **III** | $-0.054$ | $6.198^{*}$ | 0.9 | No |
| | (0.038) | (3.709) | | |
| **IV** | $-0.185$ | $26.362^{**}$ | 0.7 | No |
| | (0.115) | (11.159) | | |
| **POST** | $-0.211$ | $39.400^{*}$ | 0.5 | No |
| | (0.198) | (20.600) | | |

**Figure 1: Historical Funding–Basis Elasticity (Scaled $\beta$), ETH perpetual.** The dashed line represents the theoretical benchmark $\beta = 1$. We plot the scaled funding–basis elasticity ($\beta_{\text{scaled}} = \beta/\text{median}(f_t/\rho_t)$) and funding rate over 20 months. The scaling adjustment is for better interpretation of dynamic funding elasticity. We see the value of funding elasticity consistently over the theoretical benchmark value during the period of key market events. The funding rate was consistently elevated from November 2020 to April 2021. A positive funding rate implies that the futures is over spot. An increase in the funding rate implies increased investors' participation. The $\beta \approx 1$ levels are consistent with an efficient pricing regime.

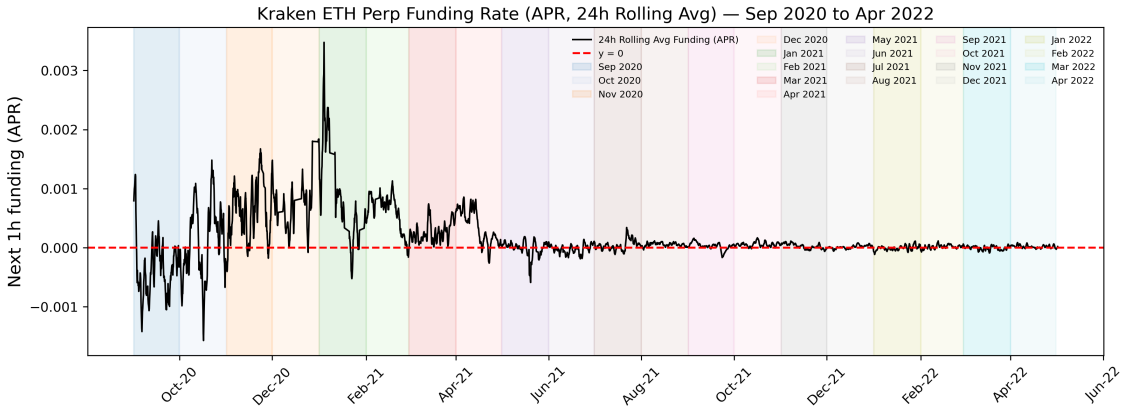


**(a)** 24-Hour Rolling Funding Rate

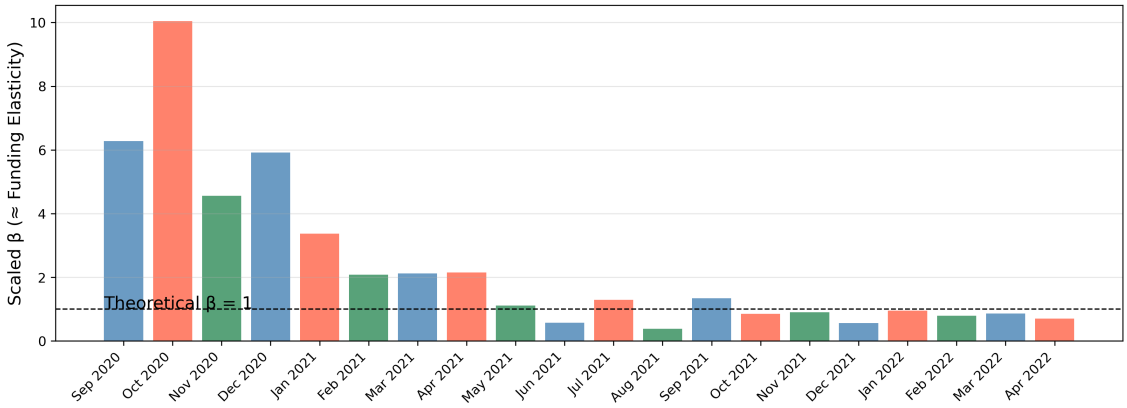


**(b)** Scaled Funding Elasticity

**Figure 2: Kraken ETH Perpetual Funding Rates Around Major Market Crashes.** Hourly annualized funding rates (APR) are shown for three crisis episodes: (a) Terra collapse (May 2022), (b) FTX collapse (Nov 2022), and (c) SVB collapse (Mar 2023). Shaded regions denote the event windows. Each episode shows a sharp deviation of funding from equilibrium levels, consistent with transient dislocations in the arbitrage mechanism. The red dotted horizontal line is the inflection point where the funding regime changes from Long to Short. Above the red line, futures trade over spot.

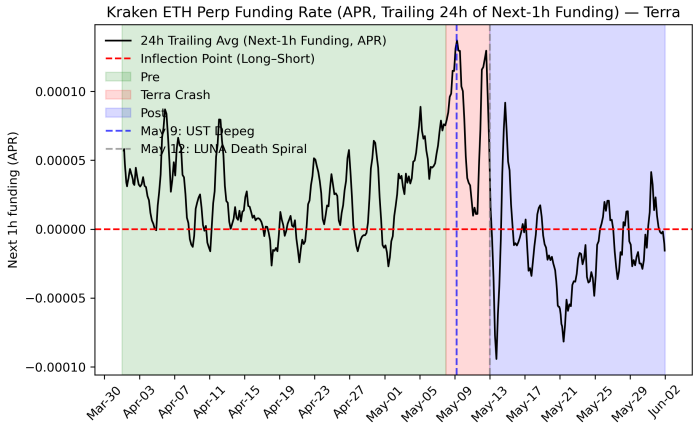


**(a)** Terra Crash Onset (May 9, 2022)

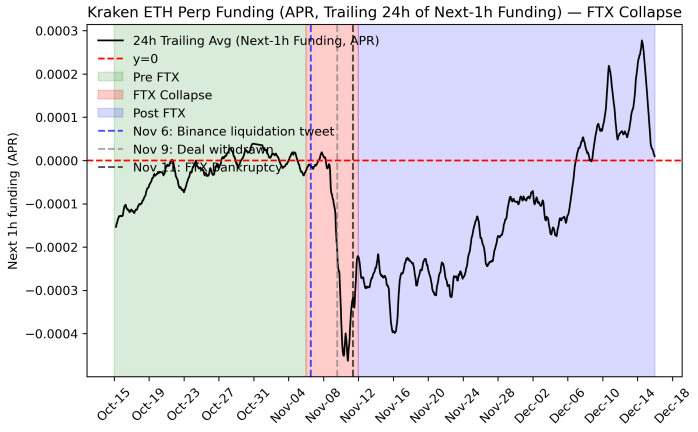


**(b)** FTX Collapse Onset (Nov 6, 2022)

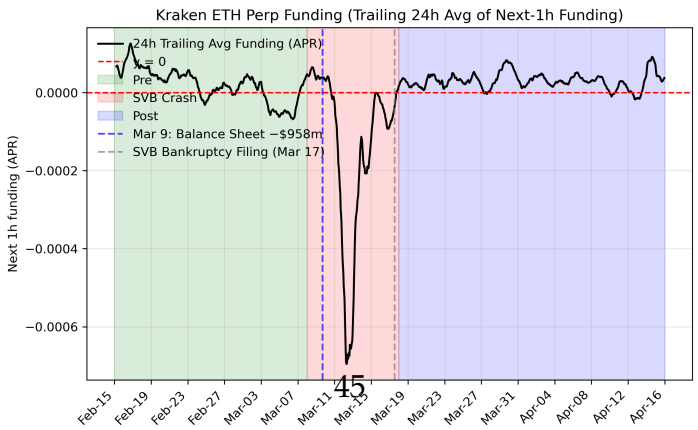


**(c)** SVB Collapse Onset (Mar 9, 2023)

**Figure 3: Funding–premium elasticity on Kraken ETH perpetuals around major stress events.**
Each panel plots next-hour funding (APR) vs. the prior-hour perp–spot premium (APR) for pre-, crash-, and post-event windows with within-window OLS fits: (a) Terra—flatter slope and higher dispersion (impaired arbitrage); (b) FTX—tight, nearly linear mapping with negative premium (shorts pay long); (c) SVB—downward level shift with slope largely preserved (funding-liquidity shock).

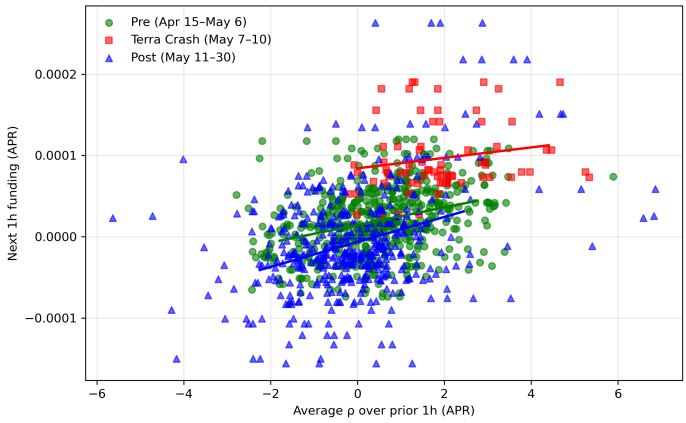


**(a)** Terra (May 9, 2022)

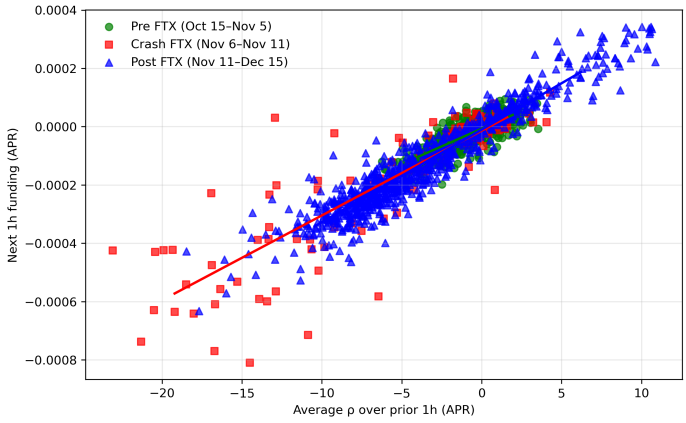


**(b)** FTX (Nov 6, 2022)

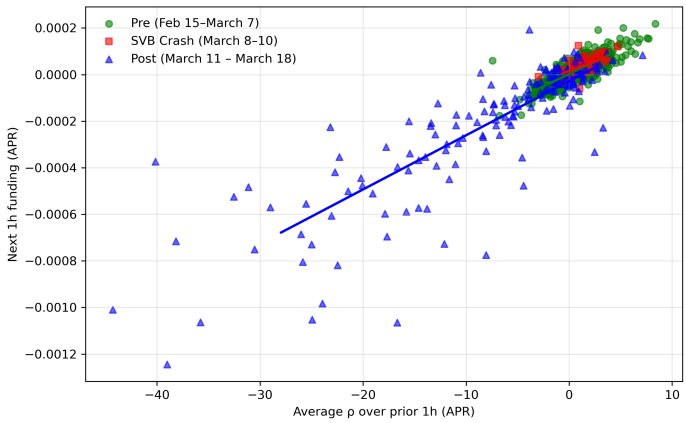


**(c)** SVB (Mar 10, 2023)

**Figure 4: Funding–Basis Elasticity Across Three Stress Episodes**
This figure reports the scaled funding–basis elasticity, $\beta_{\text{scaled}} = \beta/\text{median}(f_t/\rho_t)$, for Kraken ETH/USD perpetual swaps during three systemic events: Terra (May 2022), FTX (Nov 2022), and SVB (Mar 2023). The dashed horizontal line marks the benchmark $\beta = 1$. Elasticity collapses during the Terra run but remains above 1 during FTX and SVB, indicating preserved arbitrage function.

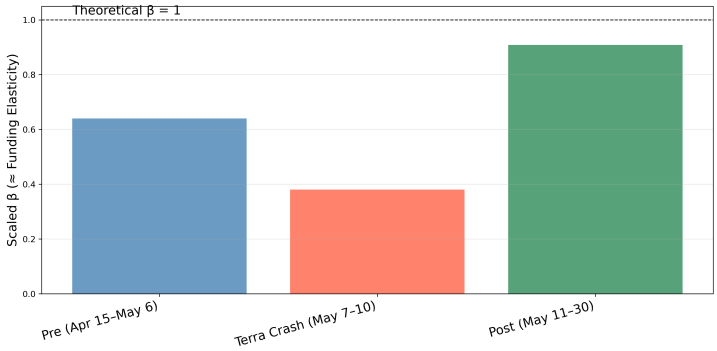


*Panel A. Terra Collapse (May 7–10, 2022)*

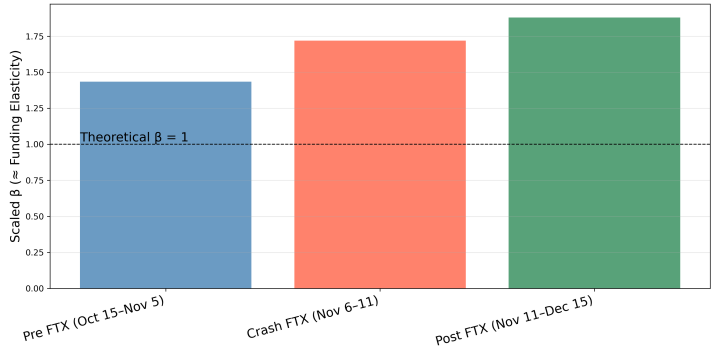


*Panel B. FTX Collapse (Nov 6–11, 2022)*

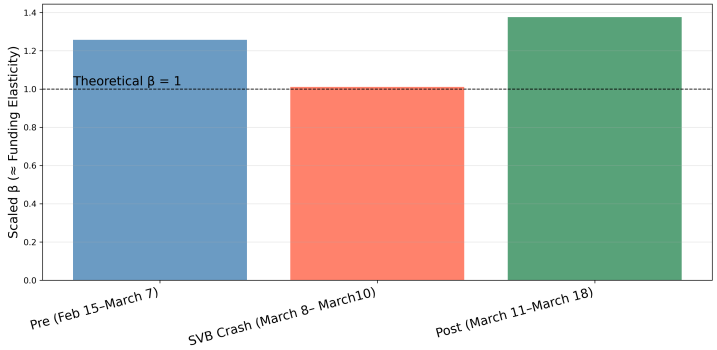


*Panel C. SVB Failure (Mar 8–10, 2023)*

**Figure 5: Volatility – Terra Period**
This figure plots hourly Ethereum (ETH) volatility ($\sigma$) in percent from April 15 to May 30, 2022, covering the Terra collapse. The series is partitioned into three event windows: Pre-Crash (blue, April 15 - May 7), Crash (red, May 7 - May 12), and Post-Crash (green, May 14 - May 30). The dashed horizontal line at 1.8% marks the volatility threshold $\sigma^*$ estimated from the global-games model (Table 7). During the Crash window, volatility spikes breached this threshold, triggering the coordination failure and run on arbitrage capital that collapsed open interest.

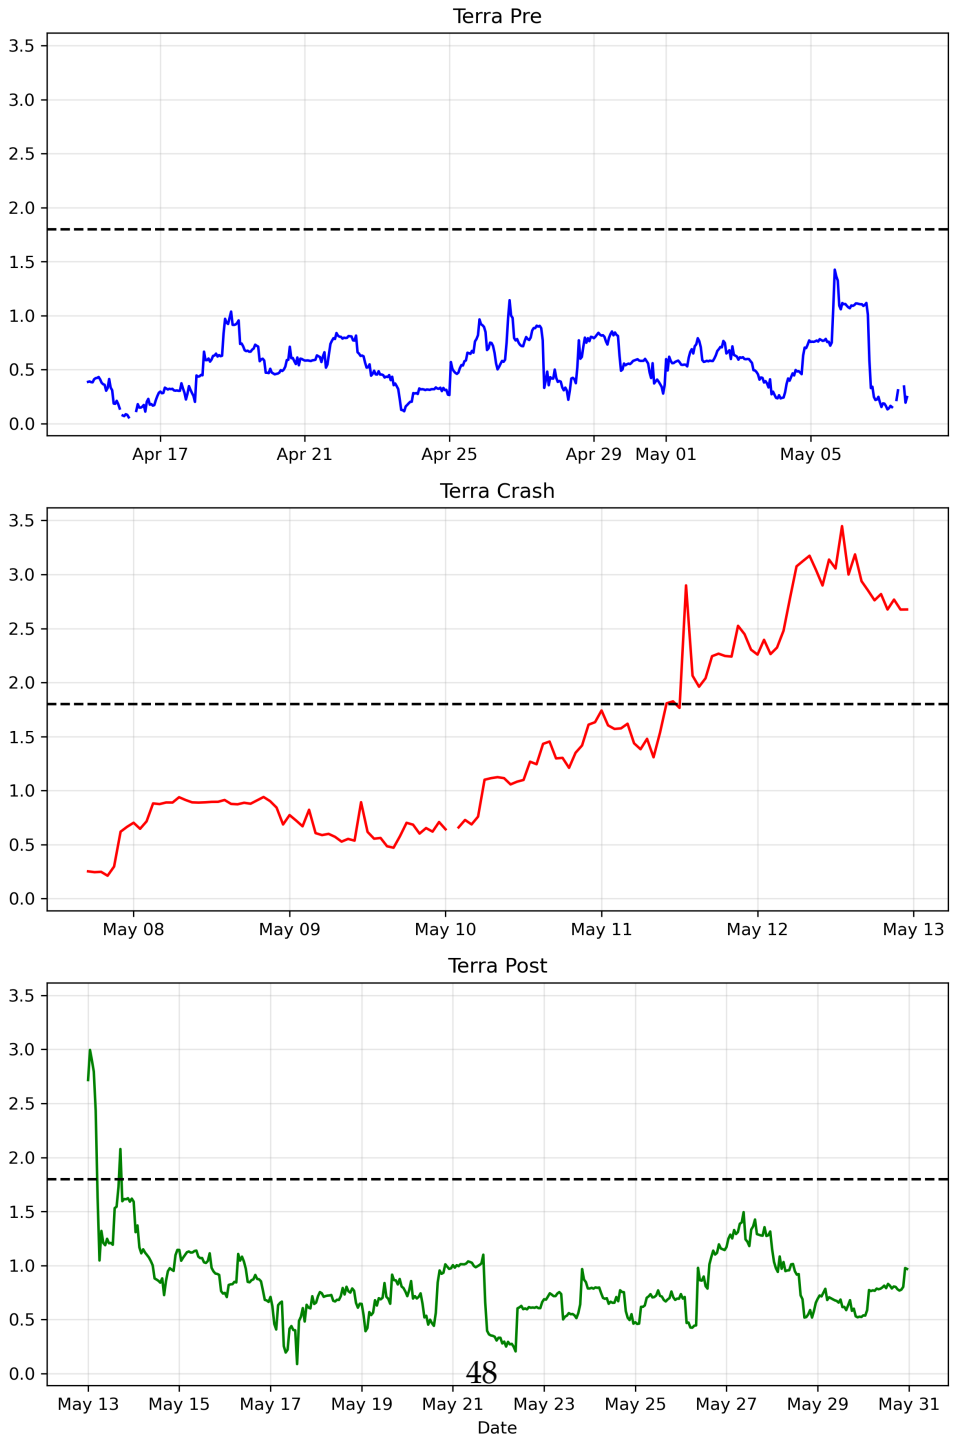

**Figure 6: Volatility – FTX Period**

This figure plots hourly Ethereum (ETH) volatility ($\sigma$) in percent from Oct 15 to Nov 30, 2022, covering the FTX collapse. The series is partitioned into three event windows: Pre-Crash (blue, Oct 15 - Nov 5), Crash (red, Nov 6 - Nov 9), and Post-Crash (green, Nov 10 - Nov 30). The dashed horizontal line at 0.9% marks the volatility threshold $\sigma^*$ estimated from the global-games model (Table 7). During the Crash window, volatility spikes breached this threshold, but the run was not triggered as $\sigma^*$ was statistically insignificant.

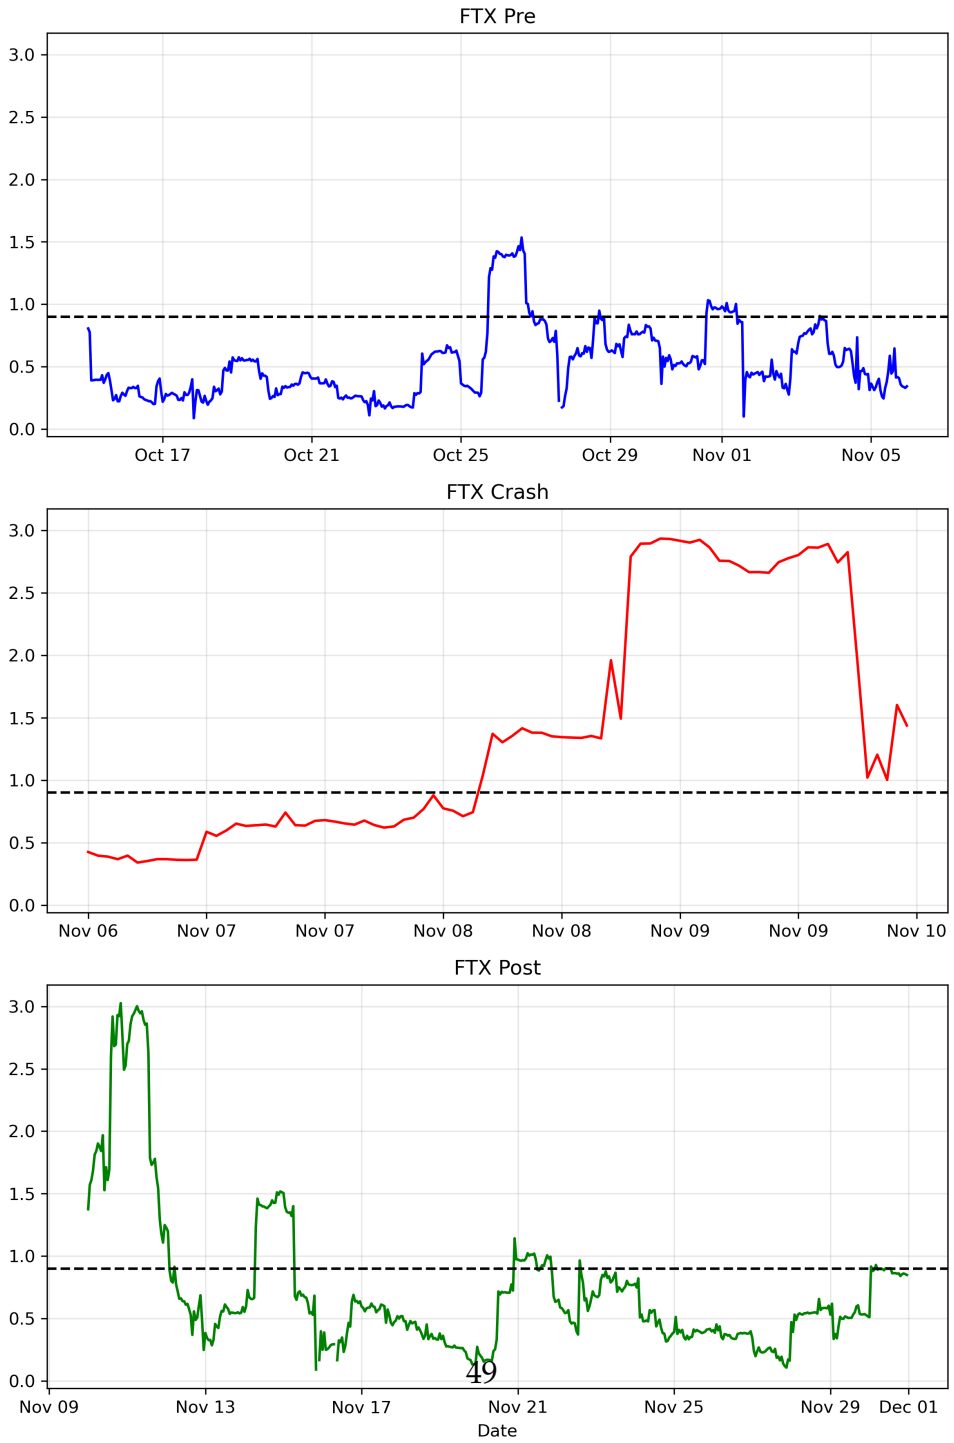

**Figure 7: Volatility – SVB Period**

This figure plots hourly Ethereum (ETH) volatility ($\sigma$) in percent from Feb 15 to March 13, 2023, covering the SVB collapse. The series is partitioned into three event windows: Pre-Crash (blue, Feb 15 - Feb 5), Crash (red, Feb 8 - Feb 9), and Post-Crash (green, Feb 10 - Feb 13). The dashed horizontal line at 1.1% marks the statistically significant volatility threshold $\sigma^*$ estimated from the global-games model (Table 7), Post Crash. Pre and during the crash window, volatility was below the threshold, and the run was not triggered. Moreover, $\sigma^*$ was statistically insignificant during these periods (0.6% and -0.1 %).

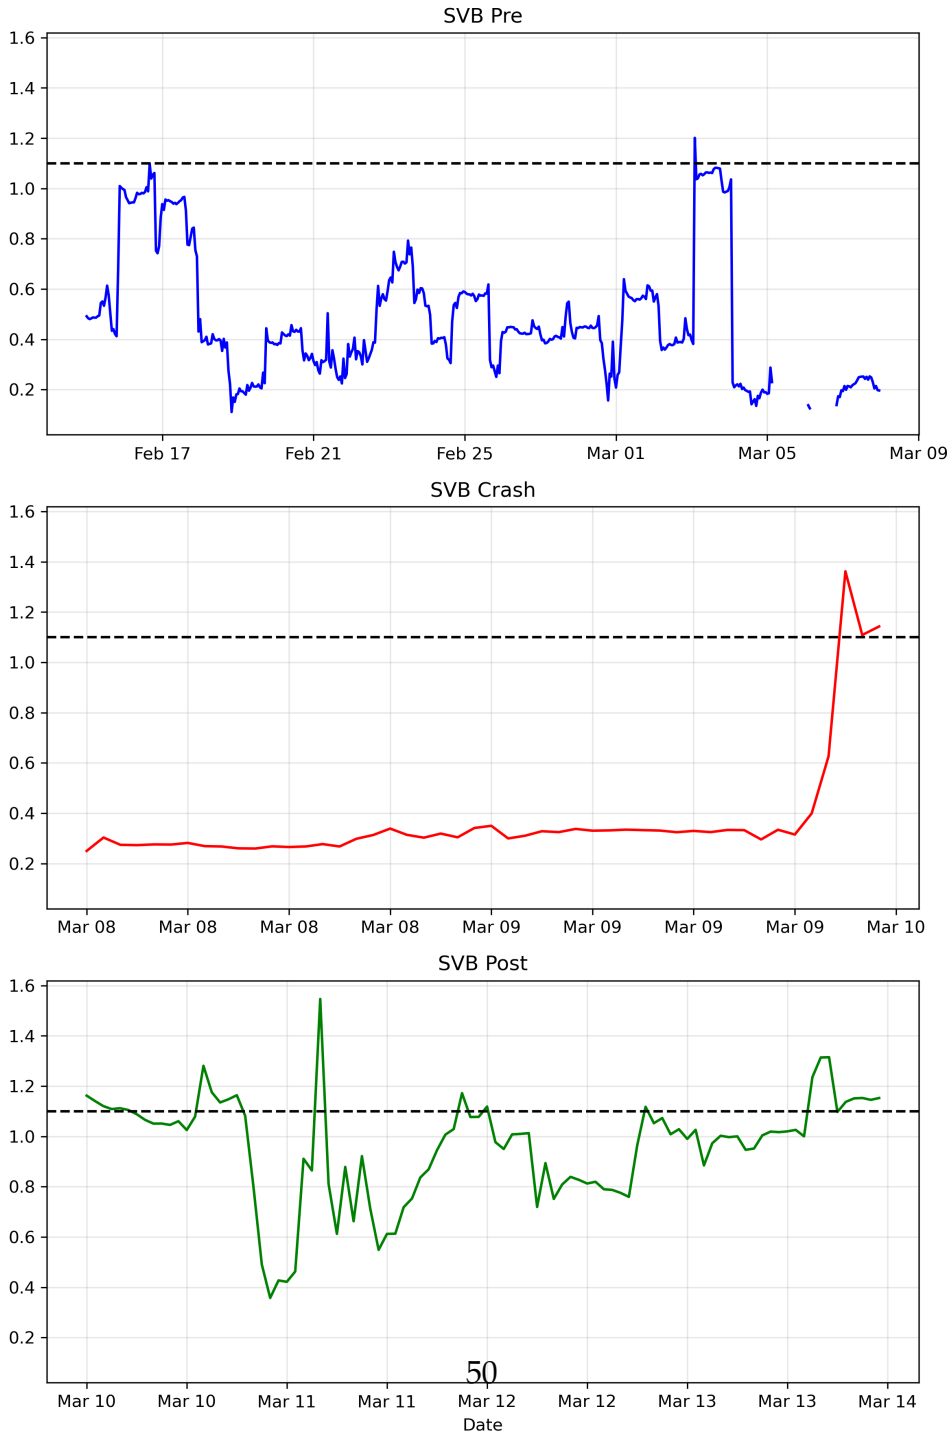

**Figure 8: Ethereum Perps Open Interest Around Major Market Crashes.**
Open Interest (notional) for three crisis episodes: (a) Terra collapse (May 2022), (b) FTX collapse (Nov 2022), and (c) SVB collapse (Mar 2023). Shaded regions denote the event windows. The Terra episode shows the steepest decline in Open Interest. The FTX episode shows an increase in notional, while the SVB crash episode shows a slight decline.

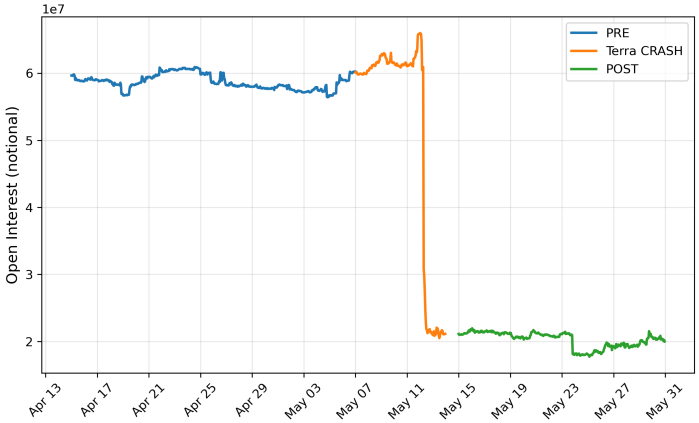


**(a)** Terra Crash Onset (May 7, 2022)


**(b)** FTX Collapse Onset (Nov 6, 2022)

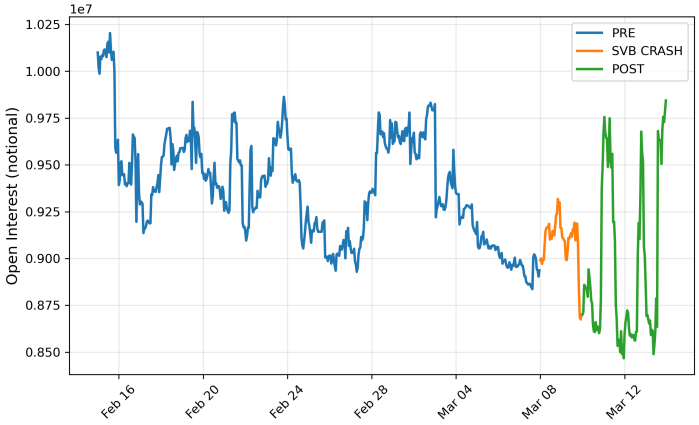


**(c)** SVB Collapse Onset (Mar 8, 2023)

**Figure 9: Ethereum Adverse Selection Costs Around Major Market Crashes.** Adverse Selection Costs (USD) are shown for three crisis episodes: (a) Terra collapse (May 2022), (b) FTX collapse (Nov 2022), and (c) SVB collapse (Mar 2023). Shaded regions denote the event windows. The Terra episode shows a sharp increase in Adverse Selection Costs, followed by the FTX and SVB crash episodes.

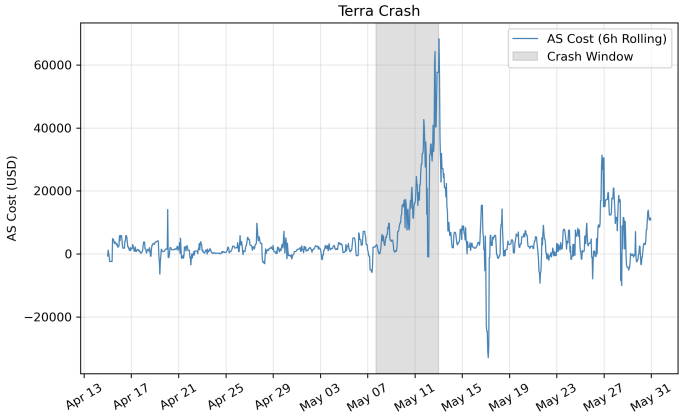


**(a)** Terra Crash Onset (May 7, 2022)

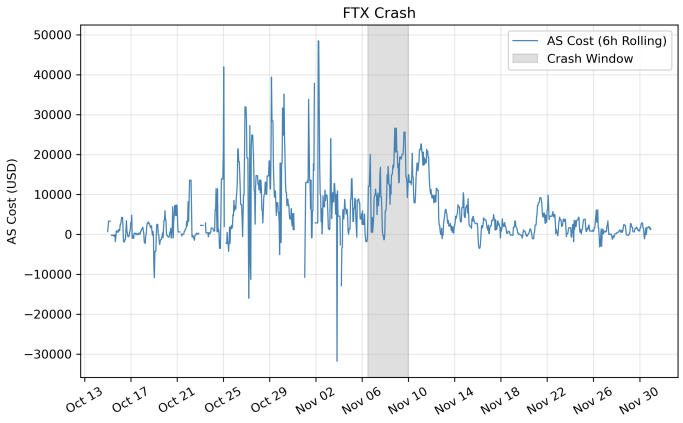


**(b)** FTX Collapse Onset (Nov 6, 2022)

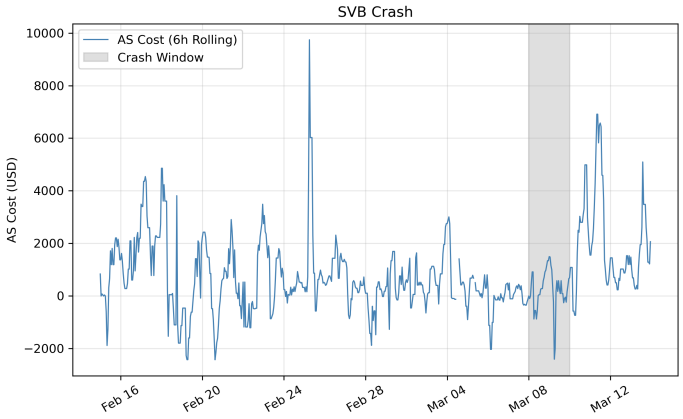


**(c)** SVB Collapse Onset (Mar 8, 2023)

**Figure 10: Scaled Funding–Basis Elasticity ($\beta$) Across China Crackdown Phases.** Each bar shows the scaled funding elasticity ($\beta_{\text{scaled}} = \beta/\text{median}(f_t/\rho_t)$) during the 2021 China cryptocurrency crackdown. The scaled elasticity $\beta_{\text{scaled}} = \beta/\text{median}(f_t/\rho_t)$ with $\text{median}(f_t/\rho_t) = 1.69 \times 10^{-5}$ yields values of 1.5, 0.09, 0.89, 0.53, 0.82, and 0.77 for the respective windows. The elasticity falls sharply to 0.09 during the high–asymmetry Phase I (May 21 –June 10), indicating impaired arbitrage and weak funding–basis linkage, and gradually recovers toward unity once the regulatory stance becomes clearer.

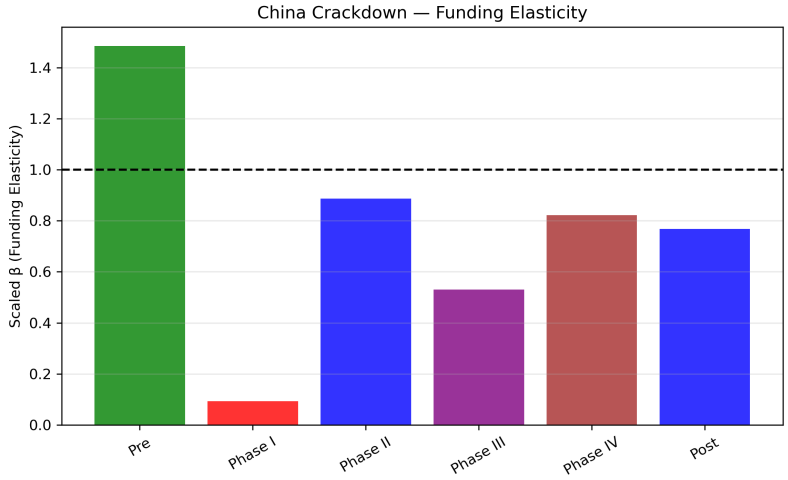

**Figure 11: Volatility – China Crypto Ban Period**
This figure plots hourly Ethereum (ETH) volatility ($\sigma$) in percentage terms from May 15 to September 24, 2021, covering the China Crypto regulatory period. The series is partitioned into four event windows. The dashed horizontal line at 4.64% during Phase I marks the statistically significant volatility threshold $\sigma^*$ estimated from the global-games model (Table 11). During other phases, $\sigma^*$ was statistically insignificant, and the run was not triggered despite volatility crossing the threshold.

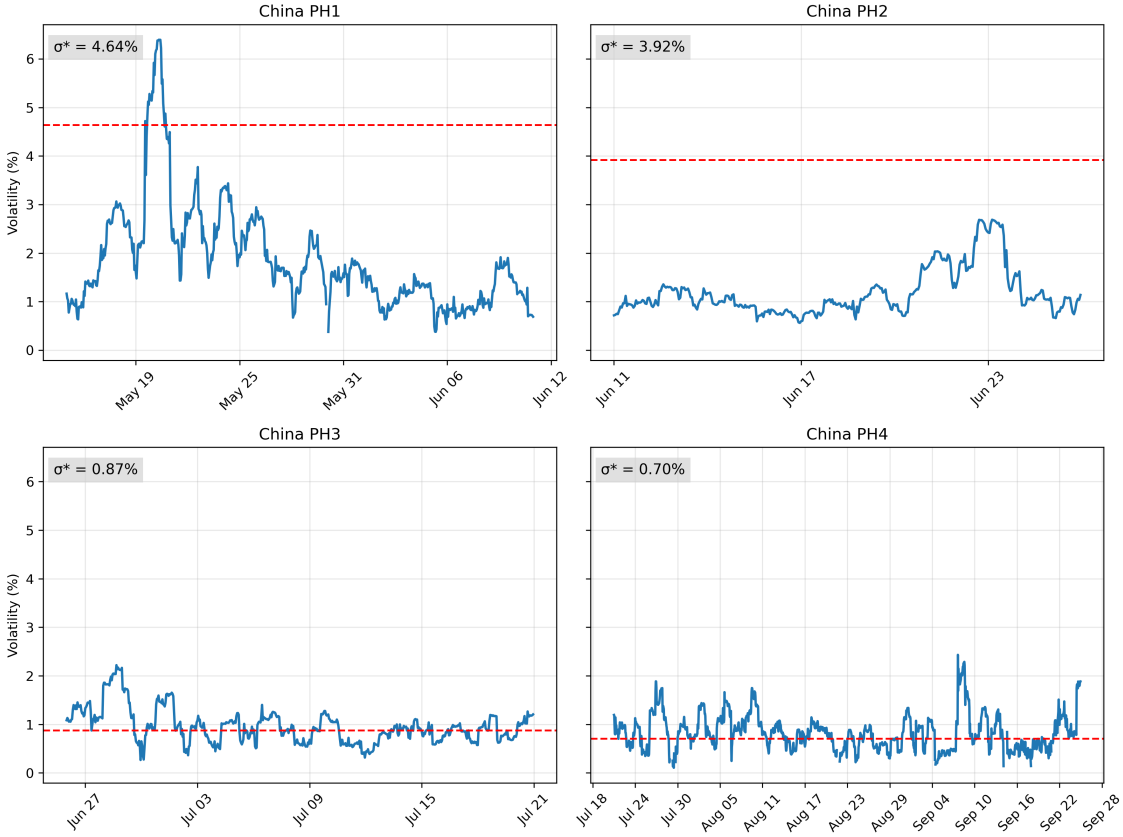

**Figure 12: Open Interest Across China Crackdown Phases.**
The figure shows the evolution of open interest during the 2021 China cryptocurrency crackdown. Open Interest falls sharply by 80% during the high–asymmetry Phase I (May 21 –June 10), indicating a run on arbitrage capital, and gradually recovers with a significant inflow of arbitrage capital during Phase IV, as adverse selection costs decline.

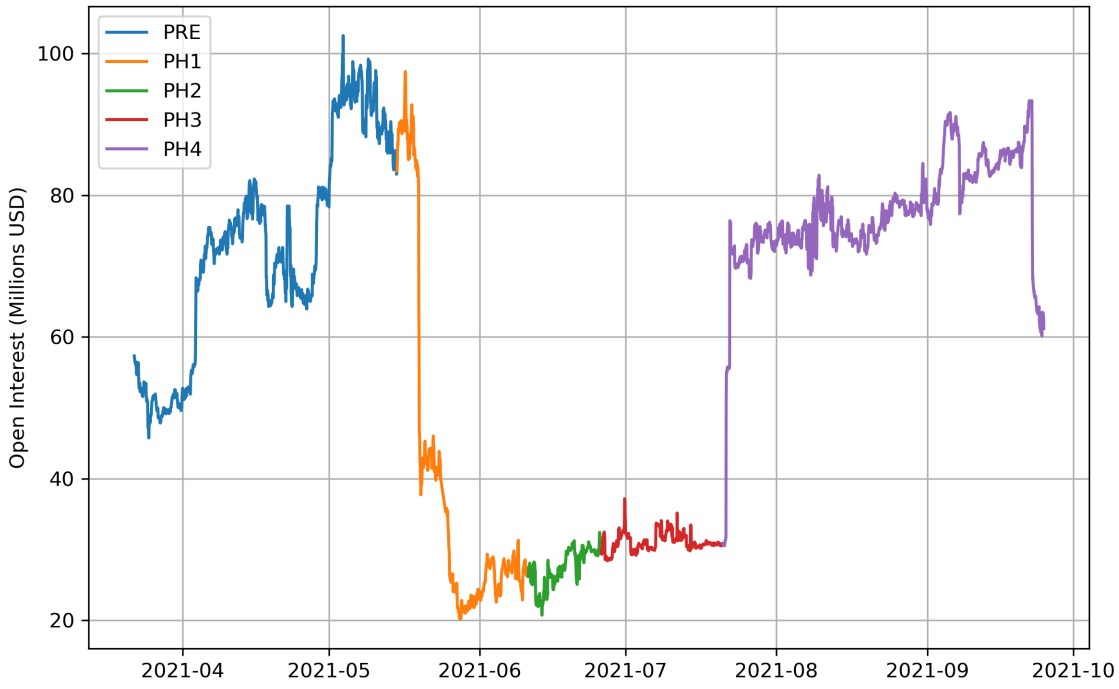

**Figure 13: Funding Dynamics and Arbitrage Mapping During China's 2021 Crackdown.** Panel (a) shows the 24-hour trailing average of next-hour ETH perpetual funding rates (APR) on Kraken, with shaded regions marking key policy phases: the 60-day pre-crackdown period (green), Phases I–IV of varying information asymmetry (orange → red → purple → brown), and the post-ban window (violet). Panel (b) plots raw 1-hour funding rates against prior-hour mark–index premia (APR) for the same windows. Funding levels turn negative and volatility increases from late May 2021 as uncertainty about mining and payment restrictions widens. The slope of the funding–basis mapping ($\beta$) flattens sharply during Phase I, indicating impaired arbitrage under high information asymmetry, and then partially recovers once the policy stance becomes symmetric and predictable (after Sept 24).

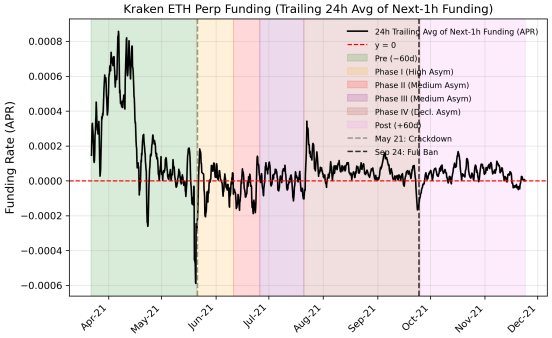


**(a)** Trailing 24h Avg Funding

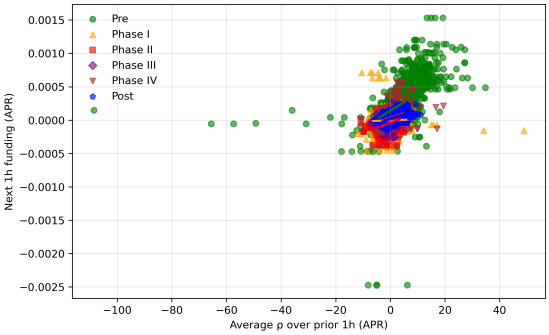


**(b)** Funding–Basis Scatter by Asymmetry

# A Perpetual Pricing Terminology – Kaiko data

## Index Price

The *index price* represents the reference spot value of ETH/USD. It is computed by Kraken as a time-weighted median of ETH/USD spot prices across major exchanges (e.g., Coinbase, Bitstamp, and Kraken Spot). The index price serves as the contract's anchor and is intended to reflect the fair market value of the underlying asset. Because it aggregates multiple venues, the index price is robust to idiosyncratic noise or manipulation on any single exchange.

## Mark Price

The *mark price* is the exchange's internal fair-value estimate of the perpetual contract. It is not the last trade price but a smoothed measure used for mark-to-market profit and loss calculations and for determining funding payments and liquidation triggers. Kraken defines:

$$\text{Mark Price}_t = \text{Index Price}_t \times (1 + \text{Premium Index}_t), \tag{18}$$

where the *Premium Index* is a rolling weighted average of the percentage difference between the perpetual's traded price and the index price over the recent interval. As a result, the mark price tracks the index price closely but adjusts for persistent deviations between the two markets.

## Funding Rate

The *funding rate* is the periodic interest-like payment exchanged between long and short positions that keeps the perpetual price anchored to its spot reference. When the perpetual trades above the index price (*positive basis*), the funding rate is positive, causing longs to pay shorts. Conversely, when the perpetual trades below the index (*negative basis*), the funding rate is negative, and shorts pay longs. The funding mechanism therefore enforces convergence between the perpetual and spot markets:

$$\text{Funding Rate}_t = \text{Interest Adjustment}_t + \text{Premium Index}_t, \tag{19}$$

$$f_t > 0 \;\Rightarrow\; \text{longs pay shorts}, \quad f_t < 0 \;\Rightarrow\; \text{skurthorts pay longs}. \tag{20}$$

## Relation to the Spot Market

The difference between the mark and index prices,

$$\rho_t = \ln(\text{Mark Price}_t) - \ln(\text{Index Price}_t),$$

is the *perpetual basis*—a measure of the deviation of the futures market from its spot reference. Arbitrageurs monitor this basis and exploit positive or negative spreads through offsetting positions in the spot and perpetual markets. The funding mechanism, by periodically adjusting cash flows between longs and shorts, serves as an endogenous force driving $\rho_t$ toward zero. In equilibrium, the mark price should closely track the index price, and funding payments adjust to enforce this parity. Temporary deviations—such as those observed during the May 2022 Terra collapse—reflect transient limits to arbitrage or liquidity constraints in the perpetual market.

# B Institutional Background

## B.1 Cryptocurrency Spot Markets

The cryptocurrency spot market, where assets are traded for immediate delivery, has evolved from a retail-dominated niche into a complex institutional-grade ecosystem. This transformation is driven by advancements in market infrastructure, regulatory clarity, and the development of sophisticated financial instruments. The core of the spot market involves the immediate exchange of a cryptocurrency, such as Bitcoin or Ethereum, for fiat currency (e.g., USD) or another digital asset. Unlike derivatives, which derive value from an underlying asset, a spot transaction confers immediate ownership. The price at which this exchange occurs is the spot price. By mid-2025, over 19,000 tokens were listed, yet market value remains highly concentrated in Bitcoin (BTC) and Ethereum (ETH). According to CoinMarketCap's Q2 2025 update, the total crypto market capitalisation was about US $4 trillion, with 24-hour spot trading volume of roughly US $98 billion[4]. CoinGecko's Q3 2025 industry report notes that the market cap continued its upward trajectory in Q3, increasing by 16.4% to end the quarter at around US $4.0 trillion [5]. Binance remains the largest spot exchange, recording roughly US $2 trillion in spot volume in Q3 2025. Binance remains the largest spot exchange, recording roughly US $2 trillion in spot volume in Q3 2025.

## B.2 Development of Cryptocurrency Derivative Markets

The first regulated Bitcoin futures began trading on Cboe on 10 December 2017 and on CME on 18 December 2017 under CFTC oversight. CME later introduced

[4] https://coinmarketcap.com/academy/article/according-to-cmc-q2-2025
[5] https://www.coingecko.com/research/publications/2025-q3-crypto-report#:~:text=1,0T%20in%202025%20Q3

cash-settled Ethereum futures on 8 February 2021 [6]. In October 2021, ProShares launched the first U.S. Bitcoin futures-based ETF (BITO), whose first-day trading volume exceeded $1 billion, signaling strong institutional demand.

By 2025, crypto derivatives account for roughly 75–85% of total crypto trading activity. TokenInsight's [7] report notes that derivatives trading on CEXs surged to US $26.0 trillion in Q3 2025, with average daily volume US $283 billion, while spot trading volume was US $4.7 trillion. Global open interest was around US $91 billion, although Coindesk reported that open interest across CEXs reached US $187 billion in August 2025. Derivatives account for the majority of crypto trading. Kaiko data show that perpetual contracts alone represented 68% of Bitcoin derivative volume. On CME, Bitcoin options open interest reached $4 billion by Q2 2025, and Ether options daily volume increased 65% year-on-year.

## B.3 Perpetual-Futures Markets

Perpetual futures (*perps*) are non-expiring contracts that use an eight-hourly *funding rate* mechanism to anchor prices to the spot market. When the perpetual price exceeds spot, longs pay shorts; when below, shorts pay longs. This continuous payment enforces convergence between the two markets.

Perpetuals, first introduced by BitMEX in 2016, now dominate crypto derivatives. CoinGecko's *State of Crypto Perpetuals Report (2025)* [8] estimates that the top ten centralized exchanges traded about $58.5 trillion in perpetual futures during 2024, doubling 2023's figure, with open interest reaching $131 billion in December 2024. Kaiko estimates that derivatives represent over 75 % of total crypto trading

[6] https://www.cmegroup.com/media-room/press-releases/2021/2/08/cme_group_announceslaunchofetherfutures.html
[7] https://tokeninsight.com/en/research/reports/crypto-exchanges-report-q3-2025
[8] https://assets.coingecko.com/reports/2025/CoinGecko-State-of-Crypto-Perpetuals-Market.pdf

activity, with perps alone making up 68 % of Bitcoin trading volume in 2025 [9]. Decentralized perpetual protocols such as dYdX, Hyperliquid, and Aster handled over $1.5 trillion of trading in 2024, with DEX perpetual volumes growing 87% quarter-on-quarter to $1.81 trillion in Q3 2025.

**Regulatory Developments.** Until 2025, perpetual contracts were offered mainly on offshore platforms. On 21 July 2025, Coinbase Derivatives launched the first U.S.-regulated nano BTC and ETH perpetual futures under the Commodity Futures Trading Commission (CFTC) self-certification. The CFTC did not object, confirming compliance with the Commodity Exchange Act, while the SEC did not challenge ETH's commodity classification [10].

### B.3.1 Market Metrics (2025)

| Metric | Value | Source |
|---|---|---|
| Total crypto market capitalization (Q3 2025) | $4.0 trillion | CoinGecko (2025) |
| Stablecoin market capitalization | $288 billion | CoinGecko (2025) |
| Spot trading volume (Q3 2025) | $5.1 trillion | CoinGecko (2025) |
| Derivatives trading volume (Q3 2025) | $26 trillion | TokenInsight (2025) |
| Perpetual futures volume (2024) | $58.5 trillion | CoinGecko (2025) |
| Perpetual open interest (Dec 2024) | $131 billion | CoinGecko (2025) |
| Perpetual DEX volume (Q3 2025) | $1.81 trillion | CoinGecko (2025) |
| Global derivatives open interest (Q3 2025) | $91–187 billion | Coindesk (2025) |
| Crypto options market monthly volume | $8.94 trillion | CoinLaw (2025) |

**Table 12:** Key market metrics for crypto spot, derivatives, and perpetual futures markets as of 2025.

The institutional landscape of crypto markets in 2025 is characterised by enormous spot and derivatives activity, evolving regulatory frameworks, and rapid innovation.

[9] https://www.kaiko.com/reports/perps-are-coming-to-america
[10] https://www.pillsburylaw.com/en/news-and-insights/cftc-perpetual-futures-btc-eth-crypto-derivatives.html

Perpetual futures have become the most popular derivative instrument due to their flexibility and leverage. Global regulators are gradually integrating crypto assets into existing market structures: the EU's MiCA and U.S. legislation, such as the GENIUS and CLARITY Acts, aim to clarify which tokens are commodities or securities and to provide regulatory oversight. The approval of spot Bitcoin and Ether ETFs in 2024 and the launch of the first U.S.-regulated perpetual futures in 2025 mark significant milestones.