Pegged... Until It's Not: Stablecoin Risk and Market Dislocation

IAQF Student Competition 2026

Abstract

We investigate cross-currency pricing dynamics between Bitcoin quoted in U.S. dollars and stablecoins (USDT, USDC) across Coinbase and Kraken, using the March 2023 Silicon Valley Bank (SVB) crisis as a natural experiment. Employing minute-level data across six trading pairs and three market regimes (pre-crisis, crisis, and post-crisis), we document that the naive cross-currency basis between BTC/USDT and BTC/USD exploded from ±5 basis points (bps) to −170 bps during the crisis, with multi-hour persistence replacing the sub-minute convergence observed under normal conditions. We decompose this basis into stablecoin peg deviation and exchange-level friction, revealing that peg risk, not venue fragmentation, drove the majority of the dislocation. USDC fell to \$0.878 (a 14.5% de-peg after basis incorporation) while USDT simultaneously traded at a premium of +54.6 bps, producing a 12% implied USDT/USDC cross-rate dislocation. Liquidity analysis shows that stablecoin-quoted markets are structurally thinner: BTC/USDC exhibits orders-of-magnitude higher price impact than BTC/USD (Amihud), a fragility that is muted in calm periods but becomes binding under stress. We document a flight-to-quality mechanism in which trading activity permanently migrated from USDC to USDT, with USDT capturing 70–87% of stablecoin volume post-crisis. We connect these findings to the GENIUS Act of 2025, arguing that mandated reserve quality, transparency, and settlement integration are directionally consistent with reduced stablecoin risk premia and would likely mitigate similar dislocations. Post-GENIUS data provides early supporting evidence of compression below pre-2023 baseline levels, though this pattern cannot be uniquely attributed to legislation alone.

I. Introduction

The emergence of stablecoins as dominant quote currencies in cryptocurrency markets has created a parallel pricing layer for digital assets. Bitcoin can now be purchased in U.S. dollars, Tether (USDT), or USD Coin (USDC) across multiple exchanges, raising fundamental questions about cross-currency parity: if these quote currencies are all nominally "worth one dollar," should Bitcoin cost the same regardless of which one is used?

In theory, covered interest rate parity and triangular arbitrage should enforce equivalence. In practice, stablecoins carry credit risk, redemption frictions, and varying degrees of regulatory oversight that can cause their dollar pegs to break. When pegs break, the cross-currency basis between BTC/USD and BTC/stablecoin widens, liquidity fragments, and arbitrage constraints under credit risk can bind sharply.

This paper provides a comprehensive empirical analysis of these dynamics, using the March 2023 SVB crisis as a natural experiment and the GENIUS Act as a regulatory counterfactual. Our contributions are fourfold. First, we decompose the cross-currency basis into peg deviation and exchange-level friction components. Second, we document a persistent post-crisis dislocation in which the basis remained negatively biased for seven or more weeks after the shock. Third, we quantify extreme liquidity asymmetry across quote currencies and a permanent flight-to-quality from USDC to USDT. Fourth, we map five channels through which the GENIUS Act would mitigate the observed failures.

Our regulatory interpretation is deliberately non-causal. Post-2025 basis compression is directionally consistent with reduced stablecoin risk premia, but it cannot be uniquely attributed to legislation; other factors such as institutional adoption and infrastructure improvements also contribute.

I.A. The SVB Crisis and Stablecoin Contagion

On March 10, 2023, the California Department of Financial Protection and Innovation closed Silicon Valley Bank (SVB) following a bank run triggered by unrealized losses on long-duration Treasury holdings. SVB's collapse was the second-largest bank failure in U.S. history, and its repercussions extended well beyond the traditional banking sector.

Circle, the issuer of USDC, held approximately \$3.3 billion of its roughly \$40 billion in reserves at SVB. When the bank failed, uncertainty over the recoverability of these deposits triggered a rapid loss of confidence in USDC's dollar peg. Over the weekend of March 10–12, USDC fell from \$1.00 to a low of \$0.878, representing a 12.2% de-peg. The de-peg was not gradual: it occurred within hours as traders raced to exit USDC positions.

Simultaneously, USDT exhibited the opposite behavior. As the stablecoin perceived to be less exposed to U.S. banking risk (Tether's reserves were not concentrated in any single U.S. bank), USDT appreciated to approximately \$1.01–\$1.02, with a crisis-period average premium of +54.6 bps. This asymmetric response, where one stablecoin's crisis became another's safe-haven premium, is a distinctive feature of the multi-stablecoin ecosystem and created extraordinary dislocations in cross-currency BTC pricing.

The crisis was resolved on March 12, when the FDIC announced that all SVB depositors (including Circle) would be made whole. USDC began recovering toward its peg, but the damage to market confidence persisted for weeks: trading volumes permanently shifted from USDC to USDT, spreads remained elevated, and the cross-currency basis exhibited a persistent negative bias long after the peg was nominally restored.

I.B. The GENIUS Act: A Regulatory Framework

The GENIUS Act (Guiding and Ensuring National Innovation for U.S. Stablecoins), signed into law in 2025 [6], establishes the first comprehensive federal framework for stablecoin regulation in the United States. The legislation was motivated in part by the very dynamics our paper documents: the SVB crisis demonstrated that unregulated or lightly-regulated stablecoins can serve as transmission channels for traditional banking shocks into cryptocurrency markets.

Key provisions include: (i) mandatory backing by high-quality liquid assets (HQLA), primarily U.S. Treasury securities and central bank reserves; (ii) a 10% cap on exposure to any single depository institution, directly addressing the concentration risk that triggered USDC's de-peg; (iii) daily reserve attestations by registered public accounting firms, designed to reduce the information vacuum associated with slow mean-reversion after stress events; (iv) clear redemption rights guaranteeing 1:1 dollar conversion on demand; and (v) pathways for integration with the Federal Reserve payment system, enabling stablecoins to function as regulated settlement instruments.

This legislation provides a direct policy counterfactual for our empirical findings: had these provisions been in effect in March 2023, the transmission channels from banking failure to stablecoin de-peg would have been substantially altered. We examine this counterfactual through five empirically-grounded channels in Section V.D.

II. Literature Review

Our analysis integrates three established strands of financial economics: limits to arbitrage, liquidity spirals, and cryptocurrency market fragmentation.

The limits-to-arbitrage framework emphasizes that mispricings can persist when intermediaries face capital or risk constraints. Shleifer and Vishny [5] show that arbitrage capital is endogenous and can withdraw exactly when apparent opportunities are largest. Our crisis-period transition from sub-minute convergence to multi-hour persistence maps directly to this mechanism.

Brunnermeier and Pedersen [1] formalize the interaction between funding constraints and market liquidity. In their framework, reduced funding capacity widens spreads, which then worsens mark-to-market risk and triggers additional liquidity withdrawal. The USDC episode we document, involving spread blowouts, dealer retrenchment, and order-flow migration, is consistent with this feedback loop. Cespa and Foucault [2] extend this logic by showing how illiquidity can propagate across related assets; our cross-currency basis evidence is consistent with that contagion channel in a multi-quote-asset environment.

Stablecoin-specific evidence in Lyons and Viswanath-Natraj [3] highlights that peg stability depends on credible redemption and arbitrage capacity, but that these mechanisms are state dependent. Our decomposition approach complements that work by separating peg deviation from exchange-level friction at high frequency during a sharp stress episode.

In digital-asset market microstructure, Makarov and Schoar [4] document persistent cross-venue spreads and show that frictional limits can sustain deviations even in highly monitored markets. We extend the interpretation to quote-currency fragmentation, where USD- and stablecoin-denominated books can diverge materially under credit-sensitive states.

III. Data

III.A. Trading Pairs and Exchanges

We collect minute-level OHLCV candle data for six Bitcoin trading pairs across two major exchanges: Coinbase (BTC/USD, BTC/USDT, BTC/USDC) and Kraken (BTC/USD, BTC/USDT, BTC/USDC). Direct stablecoin/USD prices (USDT/USD, USDC/USD) are obtained from Kraken. For Coinbase, stablecoin prices are implied from BTC pair ratios:

$$Q_t^{\text{USDT/USD}} = \frac{P_t^{\text{BTC/USD}}}{P_t^{\text{BTC/USDT}}} \tag{1}$$

Our primary empirical sample for IAQF competition inference is March 1–21, 2023 (UTC), yielding approximately 30,240 minute-level observations per pair. This window covers pre-crisis buildup, the crisis shock, and immediate post-shock stabilization. We additionally use an extended minute-level window from February 1 to May 5, 2023 for structural context and persistence analysis, and daily candles from January 2022 through January 2026 for long-run evolution and post-2025 context. This hierarchy is used to align with IAQF event-window expectations while preserving broader structural evidence. Tick-level data from Kraken provides 113,331 trades for BTC/USDC and 146,930 trades for BTC/USDT.

III.B. Coinbase Merged Order Book

A critical institutional finding shapes our analysis: Coinbase merged its BTC/USD and BTC/USDC order books in July 2022, auto-crediting USDC deposits as USD on a 1:1 basis. Consequently, Coinbase BTC/USDC prices did not reflect the USDC de-peg during our sample period. Kraken maintained separate order books, providing the true USDC-denominated price. This divergence is not a data error but an important finding about how venue architecture mediates stablecoin risk transmission.

III.C. Regime Definitions

Within the primary March 1–21 sample (UTC), we partition the data into three regimes. The *pre-crisis* period (March 1 to March 9) captures normal market conditions with stablecoins near \$1.00. The *crisis* period (March 10 to March 15) spans SVB's closure through USDC's initial recovery. The *post-crisis* period (March 16 to March 21) captures immediate normalization. For persistence diagnostics only, we extend post-crisis tracking through May 5, 2023; for structural context, we use daily data through January 2026.

IV. Methodology

IV.A. Cross-Currency Basis

We define the naive cross-currency basis as:

$$B_t^{\text{naive}} = \left(\frac{P_t^{\text{BTC/S}}}{P_t^{\text{BTC/USD}}} - 1 \right) \times 10{,}000 \text{ (bps)} \tag{2}$$

where $P_t^{\text{BTC/S}}$ is the BTC price in stablecoin S. A positive basis indicates the stablecoin trades at a discount to USD; a negative basis indicates a stablecoin premium.

IV.B. Basis Decomposition

To separate stablecoin credit risk from exchange-level frictions, we decompose:

$$B_t^{\text{adjusted}} = B_t^{\text{naive}} - D_t^{\text{peg}} \tag{3}$$

where the peg deviation component is:

$$D_t^{\text{peg}} = \left(\frac{1}{Q_t^{S/\text{USD}}} - 1 \right) \times 10{,}000 \text{ (bps)} \tag{4}$$

and $Q_t^{S/\text{USD}}$ is the stablecoin's dollar price. The adjusted basis isolates exchange-specific microstructure frictions from the mechanical effect of peg deviation.

IV.C. Implied Cross-Rates

We derive the implied USDT/USDC cross-rate, which should equal 1.00 under peg integrity:

$$X_t^{\text{USDT/USDC}} = \frac{P_t^{\text{BTC/USDC}}}{P_t^{\text{BTC/USDT}}} \tag{5}$$

Deviations from unity quantify the relative dislocation between two instruments that both purport to represent one dollar.

IV.D. Transaction Cost Thresholds

We compute round-trip arbitrage thresholds under three fee structures. Retail taker-taker: $\theta_{\text{retail}} = 92$ bps. Professional taker-maker: $\theta_{\text{pro}} = 40$ bps. Maker-optimized maker-maker: $\theta_{\text{mm}} = 29$ bps. The no-arbitrage condition holds when:

$$|B_t^{\text{naive}}| < \theta_k \quad \text{for fee tier } k \tag{6}$$

IV.E. Liquidity Metrics

We employ four complementary measures computed on 5-minute bars (5,329 aligned observations per pair on Kraken).

The Amihud illiquidity ratio captures price impact per unit of trading:

$$\text{ILLIQ}_t = \frac{|r_t|}{\text{DVOL}_t} \times 10^6 \tag{7}$$

where r_t is the return and DVOL_t is dollar volume.

The Roll spread estimator measures the effective bid-ask spread:

$$\hat{S}_{\text{Roll}} = 2\sqrt{-\text{Cov}(\Delta P_t, \Delta P_{t-1})} \tag{8}$$

The high-low spread proxy captures intra-bar trading costs:

$$\hat{S}_{\text{HL}} = \frac{P_t^{\text{high}} - P_t^{\text{low}}}{P_t^{\text{mid}}} \times 10{,}000 \text{ (bps)} \tag{9}$$

Realized volatility is computed as the 1-hour rolling standard deviation of 5-minute returns:

$$\sigma_t^{\text{RV}} = \text{Std}\left(\{r_{t-11}, r_{t-10}, \ldots, r_t\}\right) \tag{10}$$

IV.F. Execution Cost Approximation

To map spread and illiquidity changes into dollar terms, we use an illustrative reduced-form execution-cost approximation for a one-way trade of size Q_t BTC at price P_t:

$$C_t^{\text{one-way}}(Q_t) \approx Q_t P_t \left(\frac{\hat{S}_t}{10{,}000} \right) + \lambda_t Q_t P_t \tag{11}$$

where $\hat{S}_t$ is the spread proxy and λ_t is an impact component linked to Amihud illiquidity. This specification is descriptive rather than structural; it translates observed microstructure conditions into an economically interpretable cost scale.

For cross-venue comparisons, we compute the differential:

$$\Delta C_t(Q_t) = C_{t,\text{USDC}}^{\text{one-way}}(Q_t) - C_{t,\text{USD}}^{\text{one-way}}(Q_t) \tag{12}$$

which isolates the incremental execution burden associated with stablecoin-denominated routing during stress.

IV.G. Basis Persistence and Half-Life

We model the basis as an AR(1) process within each regime:

$$B_t = \mu + \rho_1 B_{t-1} + \varepsilon_t \tag{13}$$

The half-life of basis mean-reversion is:

$$t_{1/2} = \frac{-\ln(2)}{\ln(\hat{\rho}_1)} \tag{14}$$

IV.H. Order Book Depth Proxy

Lacking Level 2 order book snapshots, we utilize the number of executed trades per 5-minute interval as a proxy for market depth. Markets with deeper order books inherently generate higher execution frequencies, whereas shallower books yield fewer matches. Using tick-level data for Kraken's BTC/USDC (113,331 trades) and BTC/USDT (146,930 trades) pairs, we aggregate individually timestamped trades into 5-minute windows and compare average trade counts across regimes to evaluate structural depth differences and crisis responsiveness.

IV.I. Flight-to-Quality and Volume Migration

To assess systematic stablecoin substitution during the crisis, we construct three metrics. First, market share at 5-minute and daily frequencies:

$$\mathrm{MS}_t^S = \frac{V_t^{\mathrm{BTC}/S}}{\sum_j V_t^{\mathrm{BTC}/j}} \times 100 \tag{15}$$

Second, net order flow direction utilizes Kraken's trade-side flags to compute net buying or selling pressure. Third, the spread ratio captures liquidity divergence:

$$\mathrm{SR}_t = \frac{\hat{S}_t^{\mathrm{USDC}}}{\hat{S}_t^{\mathrm{USDT}}} \tag{16}$$

where $\mathrm{SR}_t > 1$ indicates higher relative trading costs for USDC. Collectively, market share shifts, order flow asymmetry, and spread divergence delineate the flight-to-quality mechanism.

IV.J. Statistical Methods

Cross-currency comparisons use Welch's two-sample t-tests to accommodate unequal variances. Regression coefficients are estimated with Newey-West HAC standard errors (12 lags):

$$\hat{V}(\hat{\beta}) = (X'X)^{-1}\hat{\Omega}(X'X)^{-1} \tag{17}$$

where $\hat{\Omega}$ includes autocovariance terms up to lag 12. Autocorrelation functions characterize basis persistence across regimes.

V. Results and Analysis

V.A. Cross-Currency Basis Dynamics

Normal Conditions vs. Crisis

Figure 1 presents the naive cross-currency basis for Coinbase BTC/USDT relative to BTC/USD. Pre-crisis, the basis fluctuates within ±5 bps, well inside even the tightest maker-optimized threshold of 29 bps. Active arbitrage maintains near-perfect parity, with a half-life of 0.6 minutes confirming rapid mean-reversion.

The SVB crisis shatters this equilibrium. Between March 10 and 15, the basis plunges to −170 bps as USDT appreciates relative to USD, reflecting a flight to the stablecoin perceived as safer. This exceeds even the retail arbitrage threshold by nearly a factor of two, representing a temporary weakening of arbitrage enforcement under credit risk.

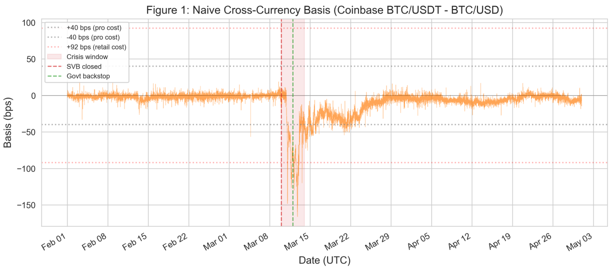

Figure 1: Naive cross-currency basis (Coinbase BTC/USDT minus BTC/USD) with transaction cost thresholds. Pre-crisis: ±5 bps. Crisis peak: −170 bps.

Transaction-Cost Band Interpretation

The fee-threshold framework offers a practical interpretation of persistence. When basis magnitudes remain inside the 29–40 bps maker/professional band, capital-efficient arbitrage can recycle quickly and inventory risk is manageable. Once basis exceeds the 92 bps retail threshold, only participants with low fees, balance-sheet capacity, and funding flexibility can continue to intermediate. During March 10–15, observed basis levels frequently moved far outside these operational bands, so even traders who identified the mispricing could not reliably monetize it without accepting material inventory and settlement risk.

This band-based view helps reconcile two facts: visible dislocations and incomplete arbitrage closure. The issue was not informational inefficiency; it was that quoted opportunities sat outside executable regions for much of the market. In that environment, cross-currency parity restoration required either a recovery in stablecoin credibility or a decline in risk constraints, not merely faster signal processing.

Basis Decomposition

Figure 2 decomposes the basis into peg deviation and exchange-level friction for both USDT and USDC. For USDT, the crisis-period basis is dominated by the peg deviation component (the USDT premium), with exchange friction contributing minimally. For USDC, the 10% de-peg generates roughly 1,000 bps of peg deviation, and exchange-level frictions compound the dislocation further. Coinbase USDC friction is near zero because its merged order book mechanically shields BTC/USDC prices from the de-peg.

This decomposition reveals that the dominant source of cross-currency dislocation is stablecoin credit risk, not venue-level microstructure frictions. Policies targeting exchange fragmentation would have been largely ineffective; the intervention needed was at the stablecoin reserve and transparency level.

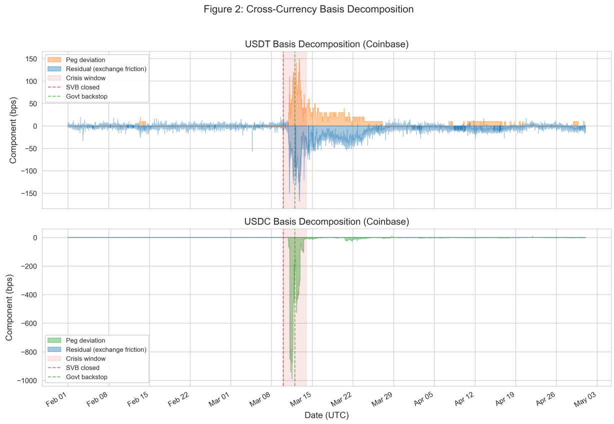

Figure 2: Basis decomposition into peg deviation (green) and exchange friction (orange) for USDT (top) and USDC (bottom).

Persistence and Post-Crisis Dislocation

Using the extended persistence window (February 1 to May 5), Figure 3 shows basis distributions by regime. The pre-crisis distribution is tightly peaked (mean = -1.3 bps, $\sigma = 3.0$ bps, n = 39,841). The crisis distribution shifts dramatically leftward (mean = -55.8 bps, $\sigma = 42.3$ bps, n = 5,020). Crucially, the post-crisis distribution does not revert to pre-crisis form: it remains wider and negatively biased (mean = -11.5 bps, $\sigma = 12.8$ bps, n = 51,513), indicating persistent post-crisis dislocation and slow mean-reversion for seven or more weeks.

Autocorrelation analysis (Figure 4) confirms this persistence. Pre-crisis, $\hat{\rho}_1 \approx 0.3$ with rapid decay, consistent with active arbitrage. During the crisis, $\hat{\rho}_1 \approx 0.99$, approaching a unit root as arbitrage capital withdraws. Post-crisis, $\hat{\rho}_1 \approx 0.93$, reflecting reduced risk appetite.

The crisis-period autoregressive coefficient ($\hat{\rho}_1 \approx 0.99$) is statistically indistinguishable from a unit root over the short crisis window. Accordingly, the implied half-life should be interpreted as a lower bound on convergence time rather than a precise estimate.

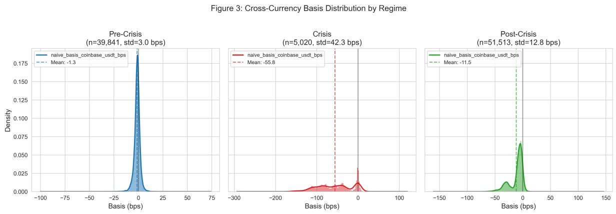

Figure 3: Basis distributions by regime. Post-crisis distribution remains wider and more negative than pre-crisis, indicating persistent post-crisis dislocation.

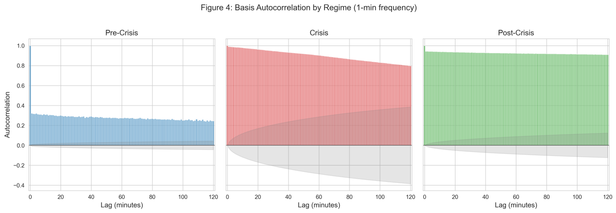

Figure 4: Autocorrelation functions by regime. Crisis-period near-unit-root behavior ($\hat{\rho}_1 \approx 0.99$) reflects temporary arbitrage withdrawal.

Volatility and Intraday Patterns

Figure 5 plots the absolute basis against BTC rolling volatility. A clear positive relationship emerges: higher BTC volatility is associated with wider basis spreads, consistent with inventory-risk models of market making. For a given volatility level, crisis observations exhibit systematically larger basis values, reflecting the additional stablecoin credit risk premium.

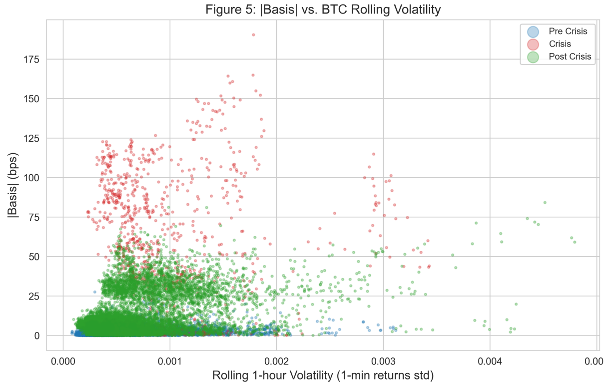

Figure 5: Absolute basis vs. BTC rolling volatility. Crisis observations are shifted upward, reflecting stablecoin credit risk beyond volatility.

Figure 6 reveals striking intraday patterns. The absolute basis peaks at approximately 75 bps during Asian trading hours (1–2 UTC) and drops to approximately 35 bps during U.S. banking hours (13–22 UTC). Pre- and post-crisis patterns are flat at roughly 2 bps and 11 bps, respectively. This demonstrates that fiat on-ramp availability during banking hours facilitates arbitrage, compressing the basis when wire transfers and ACH settlements are operational.

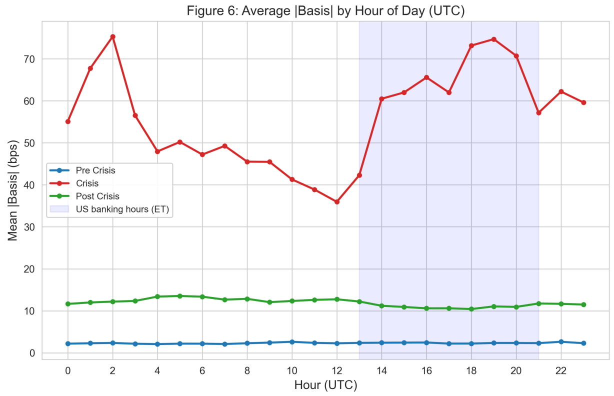

Figure 6: Average absolute basis by hour (UTC). Crisis: peaks at 75 bps during Asian hours, compresses to 35 bps during U.S. banking hours.

V.B. Stablecoin Dynamics

Asymmetric Peg Behavior and Implied Values

Figure 7 presents the implied stablecoin values derived from the no-arbitrage condition ($Q_t^{S/\text{USD}} = P_t^{\text{BTC/USD}}/P_t^{\text{BTC}/S}$) at 5-minute frequency. USDC collapsed to a minimum of \$0.878, representing a 12.2% de-peg, while USDT held steady at approximately \$0.997, confirming that the SVB shock was entirely USDC-specific. The asymmetric response explains why traders rationally fled from USDC to USDT.

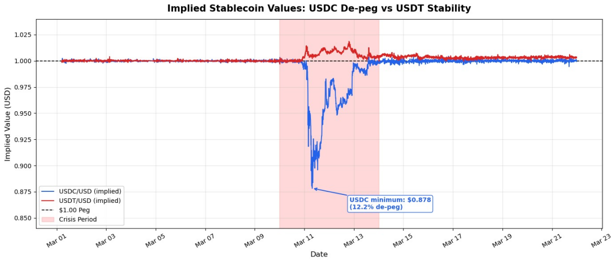

Figure 7: Implied stablecoin values: USDC/USD collapsed to \$0.878 (12.2% de-peg) while USDT/USD remained stable, confirming the shock was USDC-specific.

Figure 8 documents cross-exchange peg behavior. For USDT/USD, Kraken and Coinbase prices move in near-lockstep, confirming the USDT premium was market-wide. For USDC/USD, a dramatic divergence emerges: Kraken's price fell to \$0.878, while Coinbase's stayed near \$1.00 because of its merged order book. This constitutes a natural experiment in venue design.

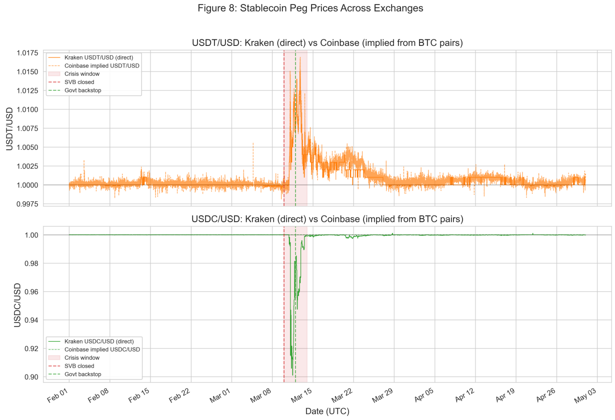

Figure 8: Cross-exchange stablecoin peg behavior. USDT/USD (top) is consistent; USDC/USD (bottom) diverges due to Coinbase's merged order book.

Regime-Specific Distribution Analysis

Figure 9 presents stablecoin deviation distributions across regimes. USDT exhibits a rightward shift during the crisis (mean = +54.6 bps), consistent with its safe-haven premium. USDC shows a massive leftward shift (mean = −242.1 bps), with the distribution spanning −1,500 to +500 bps. Post-crisis, both stablecoins exhibit residual biases: USDT at +7.5 bps and USDC at −2.7 bps.

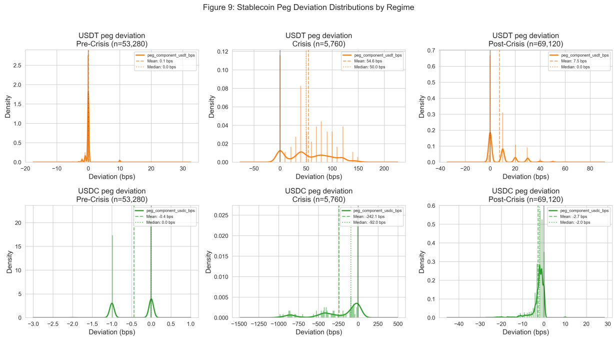

Figure 9: Stablecoin deviation distributions across regimes. USDT (top): safe-haven premium. USDC (bottom): massive de-peg during crisis.

V.C. Liquidity and Fragmentation

Structural Liquidity Asymmetry

Table 1 presents liquidity metrics across quote currencies. The volume disparity is stark: Coinbase BTC/USD processed \$9.86 billion over our sample, compared to \$207 million for Kraken BTC/USDC (a 48× gap) and \$274 million for BTC/USDT (36× less).

The Amihud ratio indicates extreme price-impact asymmetry across quote currencies: stablecoin-quoted books exhibit far higher price impact per dollar traded than USD-quoted books. Even during the crisis, when USDC volume surged 4.5×, BTC/USDC remained 2,616× more illiquid. This differential reflects both structural order book thinness and substantial baseline volume disparities. It should be interpreted as a measure of relative price impact rather than absolute fragility.

Table 1: Liquidity metrics by quote currency and regime (Kraken, 5-min bars). All cross-currency differences significant at $p < 0.001$.

Metric	USD	USDT	USDC
Pre-Crisis			
Amihud ($\times 10^6$)	1.0	2,041	13,649
Roll Spread (bps)	6.2	6.8	7.1
Trades/5-min bar	–	13.3	7.0
Volatility (%)	0.10	0.10	0.11
Crisis			
Amihud ($\times 10^6$)	1.0	–	2,616
Roll Spread (bps)	9.72	11.87	22.76
Trades/5-min bar	–	48.5	57.5
Volatility (%)	0.22	0.21	0.32
Post-Crisis			
Roll Spread (bps)	13.1	14.2	15.4
Trades/5-min bar	–	27.7	15.4

Roll spread estimates reveal that during normal conditions, effective spreads are comparable across all three quote currencies (6–7 bps). However, the crisis exposes the fragility: USDC spreads tripled to 22.76 bps ($t = 15.17$, $p < 10^{-48}$), while USD spreads rose only to 9.72 bps.

Spread Divergence

Figure 10 overlays 5-minute bid-ask spread proxies for both stablecoins. During the crisis, the average USDC spread is approximately 22 bps, while peak intrabar spikes exceed 1,400 bps; USDT remains materially lower, with peaks below 200 bps. The 22 bps average and 1,400 bps peak are different concepts: the former is a central tendency over 5-minute bars, while the latter captures tail stress in individual intervals.

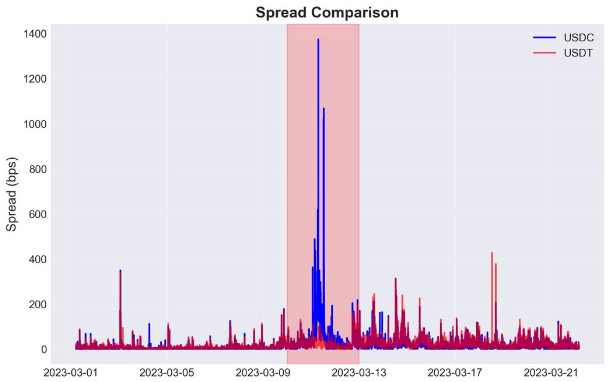

Figure 10: Average 5-minute spread (bps), with peak intrabar spikes reaching 1,400 bps.

Economic Magnitude: Stress-Test Perspective

Using the execution-cost mapping in Section IV, observed spread differentials imply economically large routing effects. For an illustrative $2.5 million BTC order (approximately 100 BTC near $25,000), the average crisis spread differential between USDC and USD books (22.76 versus 9.72 bps) implies roughly $3,260 of additional one-way quoted cost, or about $6,520 round-trip, before fees and impact. This differential is non-trivial for high-turnover strategies and compounds across repeated hedging cycles.

Tail intervals are substantially more severe. At intrabar USDC spread spikes near 1,400 bps, immediate one-way execution on the same notional would mechanically imply quoted costs approaching $350,000. While such spikes are brief and do not characterize the full crisis average, they materially alter optimal execution policy because forced liquidation risk is governed by tail liquidity, not median liquidity.

These magnitudes rationalize observed market-maker behavior. When expected adverse selection and inventory write-down risk rise, quoted depth is withdrawn even if a medium-horizon arbitrage opportunity appears attractive in mark-to-market terms. The resulting equilibrium is a temporary parity dislocation sustained by balance-sheet constraints rather than informational inefficiency.

The Liquidity Death Spiral

During the crisis, USDC markets entered a self-reinforcing liquidity withdrawal. Market makers withdrew from USDC order books to avoid devaluing stablecoin exposure. Spreads ballooned as depth evaporated. Traders responded by fleeing to USDT (net selling 995 BTC on USDC, net buying 1,180 BTC on USDT). This outflow further thinned USDC books, expanding basis persistence from under 1 minute to a multi-hour lower-bound estimate. Table 2 summarizes the dynamics.

Table 2: Basis half-life by regime, estimated from AR(1); crisis estimate shown as a lower-bound proxy due to near-unit-root dynamics.

	Pre	Crisis	Post
$\hat{\rho}_1$	0.30	0.99	0.93
Half-life (min, proxy)	0.6	> 400 (lower bound)	1.8
Mean basis (bps)	-1.3	-55.8	-11.5

Flight-to-Quality: Volume Migration

Figure 11 tracks cumulative net volume flowing from USDC toward USDT. Pre-crisis, the cumulative flow drifts mildly positive. The crisis triggers a sharp reversal around March 10–11 as panicked traders sold BTC on USDC books, temporarily driving the cumulative flow to approximately −1,500 BTC. This was followed by a steep and sustained climb to approximately +3,000 BTC, confirming that trading activity permanently migrated from USDC to USDT.

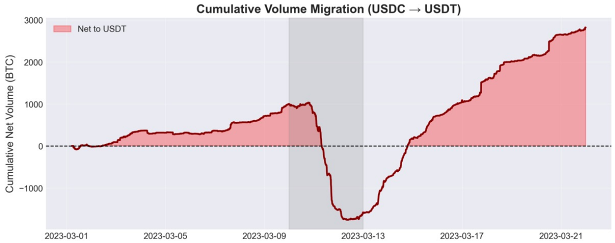

Figure 11: Cumulative volume migration (USDC → USDT). The sharp dip reflects USDC panic selling; the sustained climb to +3,000 BTC confirms permanent migration.

Figure 12 shows market share evolution at daily frequency. USDC briefly spiked to 73% on March 11 as traders rushed to exit positions, then collapsed to 18–30% for the remainder of the sample. USDT captured 70–87% of stablecoin volume post-crisis, a permanent shift that never reverted to the roughly even pre-crisis split.

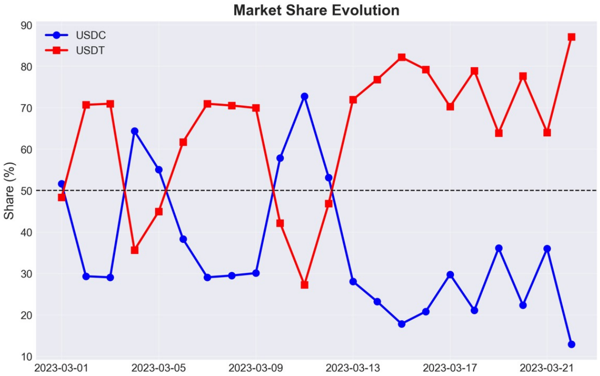

Figure 12: Market share evolution on Kraken. USDC spiked to 73% on Mar 11 (panic exits), then collapsed to 18–30%. USDT captured 70–87% permanently.

Drivers of the Spread Gap

A regression of the USDC-USD spread differential on potential drivers (Newey-West HAC, 12 lags) yields $R^2 = 0.419$ ($N = 5{,}317$):

$$\Delta S_t = \alpha + \beta_1 D_t^{\text{peg}} + \beta_2 \sigma_t^{\text{BTC}} + \beta_3 \text{VR}_t + \beta_4 \mathbb{1}_{\text{crisis}} + \varepsilon_t \quad (18)$$

The USDC de-peg is the strongest predictor ($\hat{\beta}_1 = +12.5$ bps per percentage point, $p < 0.001$). The crisis dummy adds -8.72 bps ($p = 0.035$), capturing market-maker withdrawal effects beyond the de-peg. BTC volatility is significant ($p = 0.001$) but negative, reflecting that higher volatility also brings higher volume which partially offsets illiquidity.

We acknowledge that peg deviation and quoted spreads are partially co-determined. The reported R^2 therefore reflects descriptive association rather than causal identification. The crisis dummy captures conditional market-maker withdrawal beyond the peg channel.

V.D. Regulatory Overlay: The GENIUS Act

Five Channels of Regulatory Mitigation

We identify five channels through which the GENIUS Act maps to our findings.

Channel 1: Reserve Quality. The Act mandates HQLA reserves and caps single-institution exposure at 10%. Had Circle's reserves been in T-bills rather than concentrated at SVB, a single bank failure could not have threatened the peg. The peg deviation component of Figure 2 would have been substantially reduced.

Channel 2: Transparency. Daily reserve attestations address the information vacuum associated with persistent post-crisis dislocation. The 7+ week persistence of a -11.5 bps basis (Figure 3) reflects markets lacking reliable reserve information.

Channel 3: Settlement Infrastructure. The intraday pattern in Figure 6 shows that basis compression depends on banking-hours fiat on-ramps. Integration with payment networks would enable 24/7 settlement, flattening the intraday pattern.

Channel 4: Two-Tier Ecosystem. USDT (offshore) became the safe haven while USDC (U.S.-regulated) depegged. Under a mature GENIUS Act framework, regulated stablecoins would hold diversified HQLA, inverting this dynamic.

Channel 5: Arbitrage Capital. The shift from 0.6-minute to multi-hour persistence (Table 2) reflects capital withdrawal. Regulated stablecoins reduce perceived holding risk, encouraging capital to remain deployed.

Evidence from Post-GENIUS Data

Figure 13 presents the daily absolute basis from January 2022 through January 2026. Post-GENIUS Act, the basis has compressed below the 2022–2023 pre-crisis baseline of 3–8 bps. This compression is directionally consistent with reduced stablecoin credit risk premia, deeper institutional participation, and improved confidence. The residual oscillation tracks BTC realized volatility, but the level has shifted downward from approximately 5 bps to approximately 3 bps.

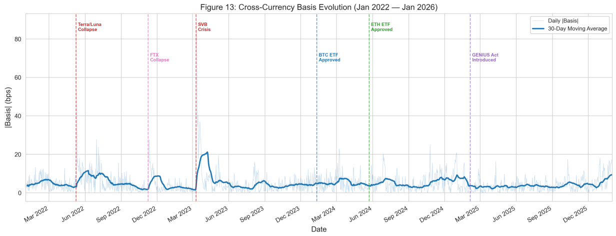

Figure 13: Full-history daily absolute basis (Jan 2022 to Jan 2026). Post-GENIUS Act basis compresses below historical baseline.

Channel Interaction and Implementation Priority

The five channels are complementary rather than additive in isolation. Reserve quality without transparent attestation can still leave a confidence gap during stress; transparency without redemption certainty may not arrest runs; and both can fail to compress intraday dislocations if settlement rails remain segmented. The crisis episode suggests that stabilization depends on a joint architecture: credible assets, verifiable disclosure, executable redemption, and continuous settlement access.

From an implementation perspective, sequence matters. The highest near-term risk reduction likely comes from balance-sheet provisions (HQLA composition and concentration limits), because those directly constrain de-peg probability. The next layer is information frequency and audit quality, which governs how quickly markets update beliefs during uncertainty. Infrastructure integration then determines how rapidly arbitrage capital can propagate that confidence into quoted prices across venues and time zones. This sequencing is consistent with our empirical ordering of dislocation channels and with the observed sensitivity of basis behavior to both perceived credit quality and payments frictions.

Policy Implications

Our findings support several recommendations: (i) reserve quality standards should prioritize short-duration government securities; (ii) real-time or daily reserve attestation is essential for rapid confidence recovery; (iii) settlement integration reduces intraday arbitrage frictions; (iv) designated market-maker obligations with spread caps (e.g., 200 bps) during stress would prevent liquidity spirals; (v) circuit breakers at 500 bps basis thresholds could pause automated trading during extreme dislocations.

While post-GENIUS compression is directionally consistent with our five-channel framework, it cannot be uniquely attributed to legislation. Institutional adoption, exchange connectivity, macro stabilization, and natural market maturation may independently contribute.

VI. Limitations

Our analysis relies on several assumptions. We use mid-price candle data, which does not capture bid-ask slippage or the true cost of execution at various order sizes. Our arbitrage profitability estimates assume unlimited recycling capital, zero borrowing costs, and perfect queue priority for the maker-optimized scenario. The convergence assumption (stablecoins revert to \$1.00) proved correct for this episode but is not guaranteed; a permanent stablecoin failure would transform apparent "arbitrage" into realized losses. Our liquidity analysis is limited to Kraken due to the Coinbase merged order book confound. The post-GENIUS Act analysis relies on early data and should be interpreted as suggestive; basis compression could reflect other concurrent factors including broader market maturation. Shifting crisis boundaries by ±2 days does not materially alter decomposition results. HAC lag specifications of 6, 12, and 24 produce qualitatively similar inferences. Finally, our single-event study design limits causal identification.

External validity is another constraint. March 2023 combines a specific banking shock, a specific stablecoin reserve configuration, and a specific exchange architecture (including Coinbase's merged USD/USDC book). Future episodes may involve different trigger mechanisms such as smart-contract failures, legal injunctions, or collateral haircuts in decentralized venues. Consequently, effect sizes reported here should be treated as regime-contingent rather than universal constants.

A related modeling limitation is omitted state dependence in arbitrage capital supply. Our reduced-form approach captures realized outcomes but does not estimate an endogenous capital-allocation function for market makers under balance-sheet and margin constraints. Developing such a structural layer would improve counterfactual precision, especially for stress scenarios where liquidity providers face simultaneous inventory, funding, and redemption uncertainty.

VII. Practical Implementation Blueprint

For market participants, the empirical results imply that stress management should be rule-based rather than discretionary. A desk-level early-warning system can combine the paper's key state variables into a single monitor:

$$\mathcal{R}_t = w_1|D_t^{\text{peg}}| + w_2|B_t^{\text{naive}}| + w_3(\text{SR}_t - 1)^+ + w_4|\text{OF}_t| \tag{19}$$

where OF_t is normalized net order flow away from the weaker stablecoin (e.g., USDC to USDT), and $(x)^+ = \max(x, 0)$. In calm regimes, $\mathcal{R}_t$ remains low and routing can prioritize fee minimization. Once $\mathcal{R}_t$ crosses a preset stress threshold, execution policy should transition automatically to preservation mode.

A preservation-mode playbook has three parts. First, execution: reduce child-order size, cut maximum participation rates, and route residual hedges to deeper USD books. Second, treasury: pre-fund fiat rails and diversify stablecoin collateral to avoid single-issuer liquidity traps. Third, governance: tighten intraday risk limits when spread-ratio and flow indicators diverge, even if mid-price volatility appears moderate. This directly addresses the core finding that normal-period spread similarity can mask severe crisis-period liquidity asymmetry.

The same framework can be extended to post-trade model validation. Risk teams can backtest whether historical VaR and execution-cost models would have remained within tolerance when the stress score rose into crisis territory. If not, model overlays should include conditional liquidity multipliers tied to peg deviation and spread divergence. In practical terms, this means treating stablecoin quote currency as an explicit state variable in both pricing and liquidation models, rather than a passive denomination choice.

A complementary control is liquidity-at-risk (LaR), which measures potential execution slippage under stressed spread and impact assumptions rather than under median conditions. For notional N_t, one practical approximation is:

$$\mathrm{LaR}_{t,\alpha} = N_t \left(\frac{\hat{S}_{t,\alpha}}{10{,}000} + \lambda_{t,\alpha} \right), \tag{20}$$

where $\hat{S}_{t,\alpha}$ and $\lambda_{t,\alpha}$ are upper-tail estimates at confidence level α. Embedding LaR directly into pre-trade checks ensures that position sizes are constrained by crisis-feasible execution rather than calm-period averages.

Operationally, this architecture should be tied to a formal escalation matrix. When stress indicators breach threshold bands, trading systems should automatically tighten participation, treasury should rebalance collateral away from the weakening quote asset, and risk should shorten permissible liquidation horizons. A documented handoff between these functions reduces discretion at the most fragile point of the cycle and improves auditability for both internal governance and external oversight.

VIII. Conclusion

This paper provides a comprehensive empirical analysis of cross-currency dynamics in cryptocurrency markets, using the March 2023 SVB crisis as a natural experiment and the GENIUS Act as a regulatory counterfactual.

Our principal findings are as follows. First, the cross-currency basis between BTC/USDT and BTC/USD is efficiently maintained within no-arbitrage bands under normal conditions (sub-minute half-life, ±5 bps), but collapses during stablecoin stress (-170 bps peak, multi-hour persistence). Second, decomposition analysis reveals that stablecoin peg deviation, not exchange-level friction, drives the vast majority of crisis-period dislocation, with USDC falling to \$0.878. Third, a persistent post-crisis dislocation remains for seven or more weeks, with the basis averaging -11.5 bps and exhibiting slow mean-reversion ($\hat{\rho}_1 = 0.93$). Fourth, stablecoin-quoted markets exhibit orders-of-magnitude higher price impact than USD-quoted markets, and this latent asymmetry becomes binding under stress, with USDC spreads reaching intrabar peaks near 1,400 bps versus 200 bps for USDT. Fifth, the flight-to-quality permanently shifted volume from USDC to USDT, with USDT capturing 70–87% of stablecoin volume post-crisis.

These findings carry direct implications for the GENIUS Act. Mandatory HQLA reserves prevent banking contagion. Daily attestations accelerate confidence recovery. Settlement integration enables 24/7 arbitrage support. Regulatory clarity inverts the perverse safe-haven dynamic. Reduced stablecoin risk encourages persistent arbitrage capital deployment. Early post-GENIUS data provides suggestive evidence of basis compression below pre-2023 baseline levels, but this compression cannot be uniquely attributed to legislation.

For trading and risk management, our results underscore that cross-currency positions in stablecoin-quoted markets carry latent risks underpriced during calm periods. The 13,649× illiquidity ratio implies that hedging or unwinding stablecoin-denominated positions during stress is far costlier than normal-period spreads suggest. Risk models should incorporate regime-dependent stablecoin peg risk and quote-currency-specific liquidity adjustments. The broader implication is that stablecoin regulation is not merely a consumer protection issue; it is a market microstructure imperative.

A key practical implication is that calm-period liquidity indicators fail to signal latent fragility. Effective spreads across USD and stablecoin pairs appear similar during normal conditions, yet underlying illiquidity asymmetry becomes binding under stress. Risk models calibrated exclusively to normal regimes will systematically underestimate crisis-period execution costs.

Finally, the March 2023 episode illustrates a broader point relevant to market design: fragility often appears first in market plumbing rather than in headline price levels. Monitoring architectures that combine peg behavior, basis dislocation, and liquidity fragmentation can identify this fragility earlier than volatility-only dashboards. As stablecoins continue to intermediate digital-asset trading, the ability to detect and price that plumbing risk will remain a defining capability for both private risk managers and public regulators.

References

[1] Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6), 2201–2238.

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

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

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

[5] Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1), 35–55.

[6] U.S. Senate (2025). GENIUS Act: Guiding and Ensuring National Innovation for U.S. Stablecoins Act. S.919, 119th Congress.