

Arbitrage Effectiveness and Stablecoin Run.

Ravi Joshi*

September 14, 2025

Abstract

Arbitrage is one of the most critical mechanisms in well-functioning financial markets. Stablecoins, designed to maintain dollar parity through arbitrage, provide a natural laboratory to study this mechanism under stress. This paper uses extremely granular data to study Terra stablecoin's arbitrage failure, which occurred 48 hours before the May 9, 2022, 5 PM depeg. I develop a generalized methodology applicable to all safe assets, using stablecoin pricing data, to measure arbitrage effectiveness in stablecoins. I further show that the declining collateral value and increasing marginal trading costs in Terra's blockchain acted as frictions to arbitrage effectiveness. I use order book data to show the microstructure of the run dynamics that followed. I show that liquidity vanished first on smaller exchanges and persisted longest on Binance, the deepest market. Results are consistent with arbitrage-run tradeoff models under extreme arbitrage concentration: unlimited participation supports price correction but amplifies run risk. My results have important implications for the stability of safe assets in general.

*rjosh15@lsu.edu

1 Introduction

Financial stability of safe assets designed to trade at par primarily depends on arbitrage mechanisms. When these mechanisms fail, persistent gaps between market and fundamental values can trigger runs. Yet despite their central importance, we have little empirical evidence on arbitrage effectiveness in real time during periods of market stress. To address this gap, I study stablecoins—an emerging class of digital safe assets marketed as dollar-equivalent instruments whose stability hinges on arbitrage mechanism. The collapse of Terra, the third-largest stablecoin with an arbitrage-based stabilization design, illustrates these risks: between May 9 and 12, 2022, it collapsed and lost \$18 billion in market value in just three days (Figure 1B).

Using Terra’s collapse as a laboratory and exploiting granular trade and order book data, I develop the first high-frequency empirical framework that identifies arbitrage failure and documents the microstructure dynamics of a run. While I apply this methodology to the stablecoin Terra, the framework generalizes to any asset class whose stability relies primarily on arbitrage, including money market funds, exchange-traded funds, and Treasury markets. Stablecoins in particular warrant attention: understanding their resilience and their implications for financial stability and monetary policy is crucial. Market capitalization surged 57% from \$150 billion to \$235 billion between November 2022 and April 2025,¹ prompting comprehensive U.S. regulation through the July 18, 2025 GENIUS Act² and positioning them as potential candidates for widespread payment use.

To see why arbitrage is central, consider the dual market structure of stablecoins. Most holders transact exclusively on secondary exchanges, where prices can deviate from the \$1 peg in response to supply–demand imbalances. In contrast, primary market transactions—direct with the issuer—occur at a fixed \$1 price. In this respect,

¹<https://www.federalreserve.gov/publications/files/financial-stability-report-20250425.pdf>

²<https://www.whitehouse.gov/fact-sheets/2025/07/fact-sheet-president-donald-j-trump-signs-genius-act-into-law/>

stablecoins resemble money market mutual funds in their primary markets while functioning more like exchange-traded funds in their secondary markets, with prices adjusting continuously to market forces. This dual market structure makes arbitrage the key mechanism that restores the peg. Traders buy in the secondary market when the price falls below \$1 and redeem in the primary market at par; conversely, when the price rises above \$1, they obtain stablecoins from the primary market at \$1 and sell them in the secondary market at the higher price. While this theoretical mechanism is well understood (Liu et al., 2023; Ma et al., 2023; Lyons and Viswanath-Natraj, 2023), the empirical evidence on its real-time effectiveness during stress periods remains limited.

The importance of the arbitrage mechanism comes into sharp focus in the May 2022 collapse of TerraUSD (UST), which exposed the vulnerabilities of stablecoins at scale. At its peak, the Terra ecosystem of UST and LUNA, the native token fundamental to arbitrage mechanism, had a combined market capitalization of about \$50 billion. Over May 9–12, 2022 (Figure 1B), their value fell to near zero, sending shock waves across digital asset markets. This paper examines two questions. First, was the arbitrage mechanism functioning effectively in the period preceding the collapse, and if not, what factors contributed to its breakdown? Second, did the subsequent market dynamics resemble those of a classic run?

I answer these questions by showing the precise timing of arbitrage failure and tracing the run dynamics across major centralized exchanges, a major venue for arbitrage activity. This paper makes three contributions. My first contribution is to test the effectiveness of Terra’s arbitrage mechanism and provide direct evidence of its failure by introducing a novel, generalizable method to identify breakdowns in arbitrage functionality. While prior studies³ (e.g., Liu et al., 2023) document the role of arbitrage trades in maintaining stablecoin pegs, I precisely characterize when and how this mechanism breaks down. I define an ‘arbitrage failure’ as a

³https://www.richmondfed.org/publications/research/economic_brief/2022/eb_22-24

structural break in the mean reversion of the peg error, where the peg error is the difference between the stablecoin’s market price and its \$1 target. Using tick-by-tick UST trade data, I track these dynamics: when the mechanism functions, arbitrage profits—measured as absolute price deviations from \$1—shrink rapidly toward zero. The framework generalizes beyond stablecoins to other safe assets, such as money market mutual funds and short-term Treasury ETFs, where redemption arbitrage stabilizes prices.

Second, I contribute to the literature on systemic risk and the market microstructure of stablecoins by providing the first microstructure evidence of the run on Terra, showing how it unfolded on centralized exchanges. The evidence indicates that the run was *global* rather than venue-specific—an insight with implications for theoretical models of coordination failure and for regulatory frameworks concerned with cross-market stability.

Third, I provide empirical support for the theoretical arbitrage–run tradeoff model of [Ma et al. \(2023\)](#) by analyzing Terra’s collapse as a real-world case at the extreme end of arbitrage concentration. Terra relied on fully decentralized, algorithmic arbitrage with effectively unlimited participation—a design that, according to the model, maximizes price stability in normal times but amplifies run risk under stress. Consistent with this prediction, I show that while UST maintained its peg for an extended period, more efficient arbitrage ultimately accelerated the run dynamics in May 2022. These findings contribute empirical grounding for theoretical models of stablecoin fragility and inform the broader policy debate on the design of arbitrage-based stabilization mechanisms.

My analysis draws on a novel dataset combining trade and order book information across stablecoins and other cryptocurrencies. Trading occurs in both primary and secondary markets: in the primary market, redemption at \$1 anchors arbitrage profits, while in the secondary market prices fluctuate with supply and demand. To study arbitrage effectiveness, I focus on the UST/USDT pair on

Binance, the most liquid stablecoin market across all exchanges during this period. This allows me to capture precise pricing of arbitrage profits and the convergence dynamics of arbitrage opportunities. To capture run dynamics during stress, I rely on high-frequency order book data for UST pairs across multiple exchanges and, to benchmark overall liquidity, I construct a synthetic order book using major stablecoin and cryptocurrency pairs (USDT/USD, USDC/USD, BTC/USDT, ETH/USDC). As a robustness check, I also examine the USDT/USDC pair—the largest stablecoin by market capitalization—which briefly depegged between May 12 and 14, 2022.

I show that Terra’s arbitrage mechanism broke down 48 hours before its collapse on May 9, 2022, as deviations from the peg ceased to be corrected despite arbitrage incentives. Arbitrage effectiveness, measured by the hourly persistence coefficient (β) from an AR(1) regression of peg deviations on their lagged values, deteriorated sharply before the collapse. While β was stable at about 0.6 through May 6, consistent with efficient peg correction, it spiked to 1 on May 7 around noon, indicating that arbitrage ceased to restore the peg. Importantly, arbitrage profits remained positive, suggesting that the breakdown stemmed from frictions in secondary markets rather than a lack of incentives.

To examine the causes of arbitrage breakdown, I primarily focus on two frictions: marginal trading costs on the Terra blockchain and the decreasing price of LUNA, the fundamental to the arbitrage mechanism. Arbitrage fails when frictions overwhelm participation and fundamentals, making price convergence impossible. On May 7, increase in trading activity pushed toward restoring the peg, but rising execution costs and LUNA’s sharp decline acted as frictions to offset these efforts, leaving deviations uncorrected. Afterward, fundamentals dominated: the collapse of LUNA made price to depeg further from \$ 1 regardless of trading activity. Overall, the results suggest that the mechanism failed not from a lack of incentives or participation, but because market frictions and collapsing

fundamentals overwhelmed arbitrage activity, preventing price correction and making convergence to the peg impossible. These findings are consistent with the framework of (Ma et al., 2023), in which a reduced set of arbitrageurs leads to persistent price deviations and peg inconsistency.

Next, I provide microstructure evidence that the run ensued after a huge withdrawal from Terra’s high-yield savings account, Anchor. It offered a fixed 19.5% APY and attracted substantial inflows, making it the dominant capital sink in the Terra ecosystem. Deposits rose from 2 billion UST in January 2022 to over 11 billion by early May 2022. On May 7, a withdrawal of 1 billion UST—approximately 9% of total deposits—occurred. A portion of these funds flowed to centralized exchanges (e.g., Binance, Coinbase), while the remainder moved to decentralized exchanges (e.g., Curve). Within hours of this withdrawal, the arbitrage mechanism broke down, and buy-side liquidity across major centralized exchanges declined sharply, leaving the order book dominated by sell-side pressure that the remaining buyers could not absorb.

To assess the impact of Anchor’s UST outflow to CEXs, I simulate standardized hypothetical trades through the full order book every minute across major centralized exchanges (CEXs). The methodology computes the volume weighted execution price (VWAP) for varying trade sizes by sequentially filling against available depth, with missing values signaling illiquidity. This provides a consistent measure of price impact and minute-by-minute liquidity. The results show that buy-side liquidity evaporated—order books thinned dramatically, and the market’s capacity to absorb sell orders vanished. During the normal period of May 1–6, the Binance order book for the UST/USDT pair was relatively balanced, with average bid depth of about \$6.4 million and ask depth of \$6.6 million. On May 7, this balance deteriorated sharply: bid depth collapsed to \$3.7 million while ask depth surged to \$10.7 million, reflecting heavy sell-side pressure. This dynamic constrained major arbitrage execution route, (CEX): as buyers vanished from centralized exchanges,

the main arbitrage channel to correct the price was blocked. Terra's price dropped because of the liquidity shock but could not be corrected because of arbitrage execution constraints. Terra could not be routed back to primary exchange for \$ 1 guaranteed redemption causing the price to further drift away from the peg. and the run ensued.

My analysis of the stablecoin collapse belongs broadly to the growing literature on digital currencies. While prior work ([Liu et al., 2023](#); [Gorton and Zhang, 2023](#); [Uhlig, 2022](#); [Catalini and de Gortari, 2021](#); [Lyons and Viswanath-Natraj, 2023](#)) emphasizes that arbitrage keeps stablecoin prices close to the peg, we know little about how to quantify the efficiency of this mechanism or identify conditions under which it fails. This paper fills that gap by introducing an empirical framework that measures arbitrage effectiveness and detects the structural break when UST's price deviations ceased to mean-revert. The analysis not only advances the stablecoin literature but also extends to other settings where the stability of publicly traded safe assets depends critically on arbitrage.

The fact that inefficient arbitrage decreases price efficiency has been shown in seminal papers by ([Shleifer and Vishny, 1997](#)) and ([Gromb and Vayanos, 2002](#)). ([Ma et al., 2023](#)) extend [Shleifer and Vishny \(1997\)](#) framework to stablecoin systems. My work builds on [Ma et al. \(2023\)](#), who show that the tradeoff between stablecoins' price stability and run risk is determined by arbitrage efficiency: concentrating arbitrage among a small set of designated players weakens peg stability but mitigates the risk of a panic run, while opening arbitrage to all investors strengthens peg stability but increases run vulnerability. I test their framework in a real-time setting with formally unlimited arbitrage participation—the extreme boundary of their model—and provide empirical validation to their theory. [Liu et al. \(2023\)](#) provide comprehensive analysis of runs on algorithmic stablecoins during the Terra-Luna crash in 2022. My work complements them by providing microstructure evidence of the run on centralized exchanges: confirming the

broad narrative but adding finer details on the run. Their work provides the first wallet-level reconstruction of the May 2022 Terra collapse. Using on-chain flows, they classify users by sophistication and show that informed wallets exited the system ahead of the crowd, the run was driven by growing sustainability concerns rather than manipulation. They show that wealthier participants exited early with smaller losses, while less informed users, suffered larger losses. While their focus is on who ran and when, I provide microscopic evidence of the run, showing liquidity dry up and price impact across centralized exchanges following Anchor's liquidity shock.

My analysis of Terra's run also builds on a large literature on panic runs and liquidity transformation (e.g., [Goldstein and Pauzner \(2005\)](#)). They show that when fundamentals cross an endogenous threshold, demand-deposit contracts switch from a unique "good" equilibrium to a self-fulfilling run equilibrium. Complementary to their findings, I document a setting in which the (positive) arbitrage spread exists, yet a run is triggered because fundamentals deteriorate.

Broadly, several other papers have explored risks associated with stablecoins. ([d'Avernas et al., 2022](#)) provide a framework to analyze how price stability can be maintained depending on the issuer's commitment to stablecoin supply. ([Anadu et al., 2023](#)) show that investors shift from riskier to safer stablecoins during periods of stress similar to the flight-to-safety behavior of MMF investors. ([Gorton and Zhang, 2023](#)) delve into the systemic vulnerabilities of stablecoins—drawing parallels to historical instances of private money issuance—and propose policy solutions grounded in monetary history and regulation. The microstructure literature under stress condition [Capponi et al. \(2024\)](#) analyze price discovery mechanisms in cryptocurrency markets, focusing on how decentralized exchange (DEX) trading activity compares to centralized exchange (CEX) liquidity provision. Complementary to these papers, I focus on stablecoin fragility under stress.

The rest of the paper proceeds as follows. Section 2 describes institutional details of Terra Ecosystem. Section 3 explains the data we use. Section 4 presents empirical findings for arbitrage failure. Section 5 shows the run dynamics. Section 6 adds the robustness results. Section 7 concludes.

2 Institutional Details

The Terra network was created in 2018 by Do Kwon and Daniel Shin through Terraform Labs. The company raised \$32 million in seed funding and another \$62 million in a 2019 token sale, offering LUNA at \$0.80. LUNA later rose to a peak price of \$119 in 2022. The total supply of 1 billion LUNA was mainly held by Terraform Labs, with early investors receiving 188 million. These funds were used to reward validators and support development. In January 2022, Terraform Labs created the Luna Foundation Guard (LFG) to help maintain the stability of the UST peg by deploying reserves during periods of market stress. LFG raised about \$1 billion from LUNA sales and, by May 6, 2022, held reserves of roughly 80,300 BTC, \$26 million in USDT, and \$24 million in USDC. During the May 2022 depeg event, LFG coordinated with Terraform Labs (TFL) and third-party trading firms to mobilize its stablecoin and Bitcoin reserves in an effort to restore the peg. A forensic audit by J.S. Held details the timing and magnitude of these interventions.⁴

The Terra blockchain was designed as a smart-contract platform with features similar to Ethereum but at a smaller scale. It supported decentralized financial applications, including borrowing, lending, savings, and trading protocols. Starting in March 2020, Terraform Labs issued an algorithmic stablecoin, **TerraUSD (UST)**, designed to maintain a fixed value of one U.S. dollar. The peg was enforced through a dual-token system with a native token, **LUNA** fundamental to arbitrage mechanism. At any time, market participants could exchange one UST for exactly

⁴See J.S. Held’s audit report: <https://lfg.org/audit/LFG-Audit-2022-11-14.pdf>

\$1 worth of LUNA, or vice versa, in the primary market. This redemption rule created arbitrage opportunities: if UST traded below \$1, arbitrageurs could buy UST cheaply, redeem it for \$1 worth of LUNA, and profit; if UST traded above \$1, they could mint(issue) UST by burning (destroying) \$1 of LUNA and sell it at a premium. This mint–burn mechanism was the core stabilization design and impacted prices in centralized exchanges where price discovery happens: UST effectively functioned as short-term debt, while LUNA represented the equity tranche absorbing losses. In this sense, the stability of UST depended directly on market confidence in LUNA’s value. This architecture makes Terra a useful setting to study the arbitrage-driven algorithmic stablecoins, and its collapse provides a clean case of systemic stress within a large but self-contained ecosystem.

2.1 Anchor Protocol

Anchor Protocol was the most important component of the Terra ecosystem. Built on the Terra blockchain, it functioned like a bank: users deposited UST and received aUST, a liquid token accruing interest. Yields—funded by collateralized borrowing in staking assets—reached up to 19.5% APY, far above both traditional benchmarks. This stable yield attracted massive inflows, effectively anchoring demand for UST and driving rapid supply growth⁵. At its peak, Anchor held over \$11 billion in UST deposits, accounting for more than 70% of all circulating UST. This extreme concentration made Anchor not just a yield platform, but effectively the *monetary base* of the Terra ecosystem. Any disruption in Anchor’s stability directly translates into instability of UST itself.

Figure 1 illustrates the interaction among the main entities in the Terra ecosystem:

⁵See Anchor Protocol documentation at <https://docs.anchorprotocol.com>

- **Terraform Labs (TFL):** The issuer of UST and driver of the arbitrage mechanism.
- **Arbitrageurs:** Entities that maintain the UST peg by exploiting price deviations between the primary market and secondary exchanges. The company (Terraform Labs, TFL) issuing UST acts as the primary market.
- **Secondary Exchanges (CEXs and DEXs):** Trading venues that facilitate liquidity provision for UST and LUNA. Price discovery happens in CEXs.
- **Anchor Protocol:** A decentralized application offering fixed yield on UST deposits, and functioning as Terra’s high-yield bank.

The vertical arrows in Figure 1 represent UST issuance and redemption between TFL and arbitrageurs. Horizontal flows depict trading activity on secondary exchanges and optional yield-seeking deposits into Anchor. The bidirectional link between Anchor and exchanges reflects capital flows, as users move funds across protocols. While arbitrageurs could place funds in Anchor to earn yield, their primary role is to maintain the UST–LUNA peg through arbitrage.

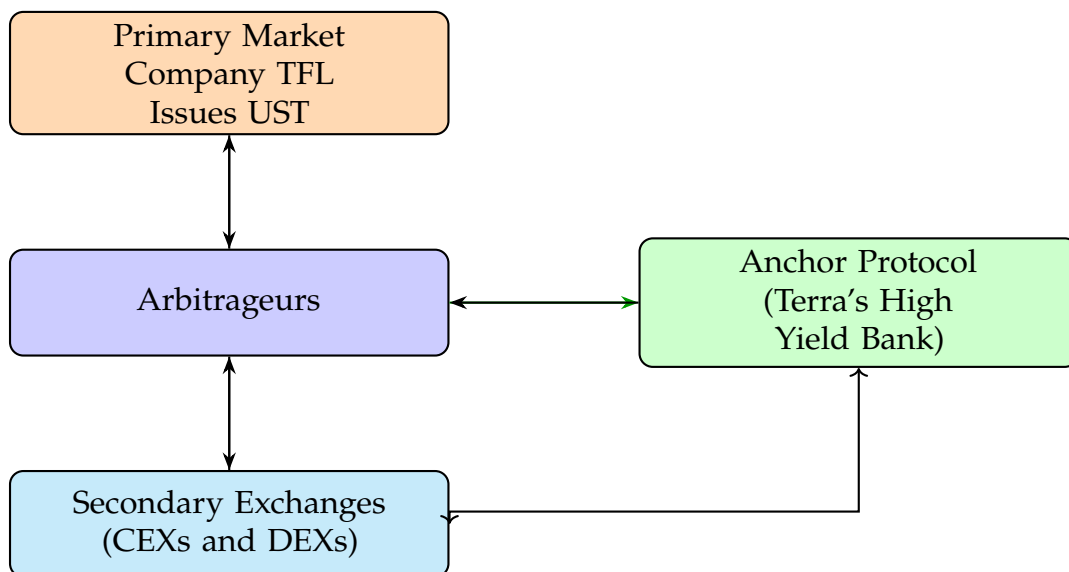


Figure 1: Terra Ecosystem

2.1.1 Lending

Anchor pooled UST deposits into a common fund, from which borrowers could draw and pay interest. The Anchor UST market contract mints (issues) the appropriate amount of aUST, and transfers it to the depositor's addresses that receive aUST tokens in exchange for their UST deposits, with aUST functioning as a transferable claim on UST, the underlying deposits held in Anchor. The value of aUST increases deterministically over time, reflecting the fixed interest rate paid to depositors. The interest on each aUST token is the difference between the aUST/UST conversion rates at the time of the deposit versus withdrawal. For example, the aUST-UST exchange rate was one at launch in March 2021, and it increased to 1.213 by March 2022. To withdraw, a depositor returned aUST to the contract, which destroyed the tokens and released the original UST plus accrued interest.

2.1.2 Borrowing

Borrowing on Anchor was over-collateralized, with users pledging bonded LUNA (bLUNA) or bonded ETH (bETH) as collateral. These bonded tokens were liquid versions of staked LUNA or ETH, designed to let investors use them immediately rather than wait through the protocol's 21-day unbonding period. Borrowers could close their positions voluntarily by "burning" bAssets, which unbonded the underlying collateral. By contrast, liquidation was automatic and involuntary: if the collateral's value fell and the loan exceeded the protocol's loan-to-value (LTV) limit—80% for bLUNA and 75% for bETH—Anchor's smart contracts seized and auctioned off the collateral to repay the debt.⁶ Both bLUNA and bETH were intended to track the value of their underlying assets and accrued staking rewards, which helped subsidize Anchor's high deposit rates. In practice, however, as

⁶See Anchor Protocol documentation at <https://docs.anchorprotocol.com/anchor-2/protocol/bonded-assets-bassets>

LUNA's price collapsed, collateral values deteriorated sharply, triggering both voluntary exits and forced liquidations.

2.1.3 Anchor Governance Token

Anchor's governance token, ANC, was designed to scale in value with protocol adoption. ANC holders could stake their tokens to create and vote on proposals, receiving a share of protocol fees in return. In this way, ANC was designed to capture part of Anchor's yield and scale in value with the size of assets under management. Protocol revenues—derived from bAsset staking rewards, excess yield, and liquidation fees—were used to repurchase ANC on the market, creating a feedback loop between Anchor's growth and token value. The total supply was capped at one billion tokens, with 40% allocated to borrower incentives, 20% to investors, 10% to the team, 15% to LUNA stakers, and 10% to the community fund. Overall, ANC was designed to reinforce Anchor's growth by rewarding borrowers, subsidizing deposit stability, and linking token value directly to the protocol's expansion.

2.2 Blockchain Infrastructure Costs

Gas fees represent blockchain infrastructure costs that directly erode arbitrage profits. When these costs spike during network congestion, they can eliminate profit margins from UST-LUNA redemption trades even when price deviations persist, thereby disrupting the peg stabilization mechanism. Gas limits, by contrast, determine the maximum computational capacity available per block and influence fee determination through supply constraints. Terra's gas limit governance differed substantially from automated adjustment mechanisms used by other blockchains. Individual validators for transactions at nodes set gas prices through node configuration parameters, while block gas limits were adjusted through

decentralized governance proposals requiring community consensus. Validators benefit from higher gas prices. During the May 2022 collapse, this governance-based system proved inadequate for rapid response to crisis conditions, with emergency interventions attempted only after arbitrage mechanisms had already failed.

3 Data

In this section, I explain the data sources.

3.1 Trade Data

I use tick-level trade records from Kaiko, a private firm that has been collecting trading information about cryptocurrencies since 2014, covering **Binance, Coinbase, Huobi, and Kraken** during May 2022. Similar data were used in prior academic studies of arbitrage in crypto markets (e.g., [Makarov and Schoar \(2020\)](#)). The sample spans UST/USDC, LUNA/USD, BTC/UST, ETH/UST, and UST/USDT pairs, capturing activity on the largest Asia- and U.S.-centric exchanges. Each record reports timestamp, price, size, and trade direction. These data allow precise measurement of peg deviations, order flow, and liquidity stress around the UST collapse. I aggregate trades to one-minute intervals to study price dislocations and the breakdown of arbitrage. I restrict my attention to the three most liquid pair: UST/USDT on Binance. I focus my analysis on the period from to May 1, 2022, to May 9, 2022. This choice is motivated by the market liquidity and arbitrage functionality during the crash period.

3.2 Order Book Data

To evaluate liquidity during stress, I use high-frequency order book data from Kaiko, covering UST pairs on major centralized exchanges (Binance, Coinbase, Kraken, among others). The sample includes UST/USD and UST/USDT pairs during May 7–9, 2022, the critical window of the depeg.

From each exchange, I retain the latest quote per minute, yielding a representative snapshot of available depth. Using these books, I simulate hypothetical market orders of varying sizes (from \$100k to \$5M) and compute the volume-weighted average execution price (VWAP). If depth is insufficient to complete a trade, the VWAP is missing, signaling illiquidity. This procedure produces a time series of estimated price impact, which I use as a proxy for execution feasibility and arbitrage costs across venues during the collapse.

3.3 Anchor Data

To examine run dynamics in TerraUSD (UST), I focus on **Anchor Protocol**, the central application in the Terra ecosystem. Anchor data are particularly informative because they track depositor behavior in real time, providing a DeFi analog to *bank withdrawals*, *asset fire sales*, and *balance sheet contractions* in traditional finance. I obtain daily balance and collateral burn and liquidation data from Flipside Crypto⁷, which records the aggregate UST deposits and withdrawals on Anchor. These balances reveal the pace and scale of investor exits during the collapse, highlighting the underlying funding fragility of the system.

⁷<https://flipsidecrypto.xyz/studio>

4 Methodology and Results

4.1 Arbitrage Functionality and Breakdown

Arbitrage was the central mechanism maintaining TerraUSD (UST) at its one-dollar target. When UST traded below \$1 in secondary markets, arbitrageurs could purchase it and redeem in the primary market for exactly \$1 worth of LUNA. Conversely, when UST traded above \$1, arbitrageurs could acquire UST at par in the primary market and sell it at a premium in secondary markets. This convertibility reduced supply when UST was undervalued and expanded supply when it was overvalued, providing a self-correcting channel that anchored the peg. Figure 2 provides a stylized illustration.

Arbitrageurs thereby linked primary and secondary markets. Their trades raised prices where UST was cheap and lowered them where it was expensive, aligning valuations across venues. Terra's design relied almost exclusively on this channel, assuming arbitrageurs would eliminate even small deviations swiftly enough to maintain dollar parity. However, this arbitrage mechanism was constrained by frictions: that arbitrageurs could act without delay, that the blockchain execution to process transactions was fast, and that markets would trust LUNA backing for redemptions. During the collapse of May 2022, these assumptions failed. Increased blockchain transaction fees slowed execution, and LUNA's price collapsed, undermining the very foundation of the peg.

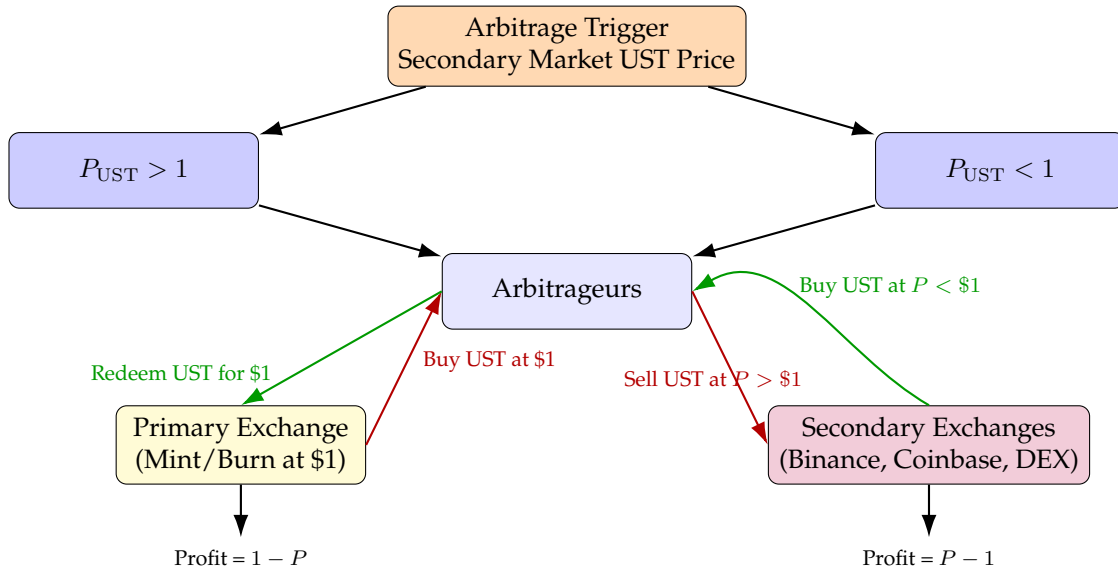


Figure 2: UST Arbitrage Mechanism: Peg Maintenance

UST and LUNA were also actively traded on large international exchanges such as Binance, KuCoin, Kraken, and Coinbase. Each platform maintained its own order book and pricing, which often diverged across venues and regions. These discrepancies created opportunities for cross-exchange arbitrage.

For instance, if UST traded at \$0.99 on Binance and \$0.995 on Coinbase, a trader could buy UST on Binance at the lower price, transfer it to Coinbase, and sell at the higher price. Similar gaps could arise within a single exchange across currency pairs or regional books. Traders acted quickly to exploit such differences, and in doing so, their trades helped align prices across markets.

Together, these arbitrage activities—both within the Terra system and across global exchanges—formed a self-correcting mechanism. As long as participants could move swiftly to capitalize on price discrepancies, their actions helped keep UST close to its one-dollar target.

4.2 Arbitrage Breakdown

The Terra–LUNA algorithm was designed to preserve UST’s peg to the U.S. dollar by ensuring arbitrage corrects prices whenever they deviated from parity. As deviations widened, the potential gains from arbitrage increased, incentivizing traders to swap UST and LUNA, and restore the peg.

Formally, the arbitrage mechanism should reduce deviations over time: when the lagged deviation $|1 - P_{\text{UST},t-1}|$ is large, arbitrageurs are expected to act, driving the current deviation $|1 - P_{\text{UST},t}|$ closer to zero. To quantify this adjustment, I estimate the following AR(1) model using tick by tick price data:

$$|1 - P_{\text{UST},t}| = \alpha + \beta|1 - P_{\text{UST},t-1}| + \epsilon_t, \quad (1)$$

where $P_{\text{UST},t}$ is the tick-by-tick Binance price of UST at time t , aggregated into hourly regressions. The coefficient β measures the persistence of peg deviations: a low β reflects effective arbitrage that quickly restores parity, while β near one signals arbitrage failure. The long-run deviation implied by the AR(1) process is given by $\frac{\alpha}{1-\beta}$, which captures the long-term arbitrage profit (LTV).

Figure 3 plots the evolution of β and $\frac{\alpha}{1-\beta}$ from May 4 to May 13. Through May 6, β remains stable around 0.6, consistent with a functioning arbitrage mechanism. On May 7, however, β spikes to one, marking a sharp breakdown in the arbitrage response. While $\frac{\alpha}{1-\beta}$ stays positive but close to zero, indicating available profit opportunities, the failure of β to fall back below unity implies that traders were unable to restore the peg despite clear incentives.

A brief decline in β on May 8 suggests some temporary price restoration, but by May 9 the coefficient again converges to one and remains there. This persistence reflects a structural failure of the stabilizing mechanism. The first clear signal of breakdown appears on May 7 at approximately 5 PM, when deviations ceased to mean-revert. Thereafter, $\frac{\alpha}{1-\beta}$ rises sharply, consistent with ever-larger arbitrage

profits that remained unexploited. Together, these dynamics reveal that after May 7 the arbitrage mechanism failed to function, leaving UST prices to drift below parity.

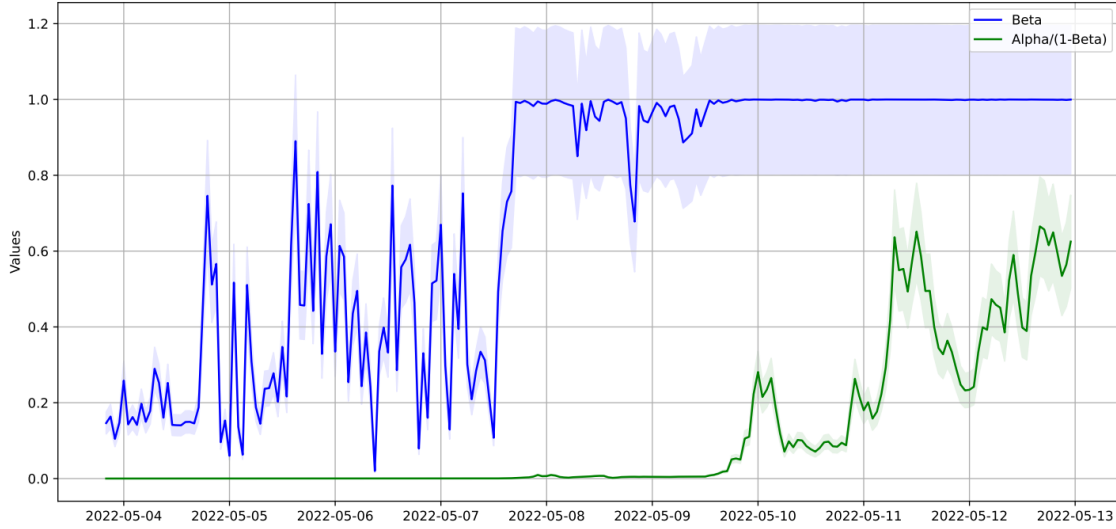


Figure 3: This figure shows the hourly values of coefficient β and the stable value of arbitrage profits in our sample from May 4, 2022 to May 12, 2022. The solid blue and green lines display, respectively, the values of coefficient β and arbitrage profits, $\frac{\alpha}{1-\beta}$, with a 95% confidence interval. A sudden increase in the value of β to 1 signifies the breakdown of arbitrage on May 7.

Figure 4 provides a detailed view of the arbitrage mechanism's performance for May 7, focusing on the impact of large withdrawals from Anchor. The plot shows the values of β and $(\frac{\alpha}{1-\beta})$ every hour, with key withdrawal events marked on the timeline. The value of β starts to consistently increase from May 7, 12 PM, and reaches a value of 1 at 5 PM. At the start of May 7, both β and $(\frac{\alpha}{1-\beta})$ exhibit stable behavior, indicating that the arbitrage mechanism is functioning effectively. A large withdrawal of 175 million USD from Anchor to DEX(Curve)([Liu et al., 2023](#)) at 12 PM is marked on the timeline, which serves as a critical point of interest. Following this large withdrawal, the arbitrage mechanism breaks down. This is evidenced by a sharp increase in β and a corresponding rise in $(\frac{\alpha}{1-\beta})$, indicating increased market risk and increased arbitrage profitability.




Figure 4: This figure shows the values of coefficient β and the stable value of arbitrage profits for May 7, 2022. The solid blue and green lines display the values of the coefficient β and arbitrage profits, $\frac{\alpha}{1-\beta}$. The vertical lines show major UST withdrawals from Anchor to CEX and DEX. A sudden increase in the value of β to 1 signifies the breakdown of arbitrage on May 7, 12 PM.

This failure is further evidenced by the daily values of β , the corresponding long-term stable value of arbitrage profit as presented in Table 1. Pre-May 7, the values of β are below 1, indicating a stable arbitrage mechanism. The model is stationary, and the arbitrage system functions as expected. On May 7, the value of β jumps to 0.999, signaling that the arbitrage mechanism is on the verge of breaking down. This sharp increase suggests that market conditions are deteriorating rapidly, and the system is approaching instability. Such a high β value implies that the market forces were insufficient to counteract the deviations from the peg, resulting in prolonged instability.

A consistent increase in β on May 7 is a critical indicator of the weakening arbitrage mechanism. It underscores the limitations of the algorithm under certain market conditions and suggests that external factors or systemic vulnerabilities may have played a role in this failure. The number of observations increases from May 7 onwards, reaching over 1.7 million by May 11. This increase suggests heightened

trading activity and market scrutiny as participants react to the unfolding arbitrage opportunities and associated risks. After May 7, β reaches and remains at 1.000, indicating a complete breakdown of the arbitrage mechanism.

Before May 7, LTV Arbitrage Profits are at 0.01 and 0.02 cents, reflecting minimal but stable arbitrage opportunities. These values suggest a consistent yet modest profit from arbitrage activities, aligning with the stable β values during this period. On May 7, LTV Arbitrage Profits show a significant jump from 0.02 cents to 0.51 cents. The rise in profits reflects the growing arbitrage opportunities due to market inefficiencies, as β approaches instability. LTV Arbitrage Profits skyrocket to 50.91 cents by May 12. This substantial increase highlights the severe market disruption as the system is now in a breakdown state with β at 1.000.

Table 1: Estimates of β and Long-Term Arbitrage Profits

Day	May 04	May 05	May 06	May 07	May 08	May 09	May 10	May 11	May 12
β	0.628	0.878	0.604	0.999	1.000	1.000	1.000	1.000	1.000
α	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$\frac{\alpha}{1-\beta}$ (cents)	0.01	0.01	0.02	0.51	0.52	10.79	16.46	45.95	50.91
Observations	16,123	19,768	15,764	41,985	77,072	190,426	751,686	1,744,320	1,566,119

Notes: This table reports daily estimates of β from the AR(1) model in equation (1), along with implied long-term arbitrage profits $\frac{\alpha}{1-\beta}$. The dependent variable is the absolute deviation from the peg, $|1 - P_{UST,t}|$. $\beta < 1$ implies mean reversion, while $\beta \rightarrow 1$ indicates arbitrage breakdown. The values of α are 0 because of rounding to three significant digits. Long-term values are rounded to two significant digits and are reported in cents. Estimates are based on tick-level UST/USDT Binance trades aggregated to hourly regressions, May 4–12, 2022.

4.3 Why did arbitrage fail?

In practice, the arbitrageur has to incur a number of transaction costs, but their magnitudes are too small to prevent arbitrageurs from executing the arbitrage trades. Most exchanges do not charge fees on a trade-by-trade basis but assign them based on the trading volume in a given month or week (Makarov and Schoar, 2020). In the absence of frictions, arbitrageurs should immediately exploit price deviations across venues, thereby restoring the peg. In practice, several frictions

can impede this process during the period of market stress. I highlight two in particular: (i) transaction costs on the blockchain and (ii) the declining value of the asset backing the stablecoin.

Transaction fees (often referred to as *gas fees* in blockchain settings) represent the cost of executing transactions on the blockchain. The blockchain has limited processing capacity and monetizes access via gas fees, effectively limiting participation during periods of high traffic. Similar to payment processing fees in traditional finance, these charges compensate validators for confirming transactions in blockchain and serve as a congestion-control mechanism. High transaction fees reduce participation because not all traders can afford to pay them. The burden of these fees is especially severe for retail participants transacting in small amounts, as the fixed costs of execution quickly erode their potential profits.

On the Terra network, every transaction—including trades—incur a gas fee, which is a small computational cost set by validators, and the fees on Terra were not fixed. They could increase when trade volumes surged. I examine hourly gas fees paid and 12-hour swap volume on the Terra blockchain, using data from Flipside Crypto. Figure 5 reports the total gas fees paid per hour on the Terra blockchain between May 1 and May 11, 2022. While these fees capture the total cost of transacting in the Terra blockchain, arbitrageurs engaging primarily with the primary market are also subject to a portion of these costs.




Figure 5: UST–LUNA on chain DEX swap volumes (left axis) and hourly gas fees (right axis) on the Terra blockchain, May 1–11, 2022. Swap volumes started increasing on May 5, with a sharp spike on May 9. Gas fees, usually near zero, began to rise on May 5 and peaked after May 9.

(Ma et al., 2023) argue that a vast majority of arbitrageurs hold only very small amounts of stablecoins and redeem in the primary market exactly as much as they purchase in the secondary market. Arbitrageur j faces quadratic trading costs: arbitraging z_j units of the stablecoin from the secondary to the primary market entails a cost

$$c_j(z_j) = \frac{z_j^2}{2\chi}, \quad (2)$$

where $\chi > 0$ captures arbitrageurs' balance sheet capacity. A higher value of χ implies lower trading costs, reflecting greater capacity to absorb arbitrage trades.

The stablecoin's secondary-market price is given by

$$p_2(\lambda) = \begin{cases} 1 - K\lambda, & \text{Issuer Solvent,} \\ \frac{1-\phi}{\lambda} - K\lambda, & \text{Issuer Insolvent,} \end{cases} \quad (3)$$

where

$$K \equiv \frac{1}{n\chi}. \quad (4)$$

Because the Terra ecosystem did not impose a cap on the number of arbitrageurs, we can theoretically take $n = \infty$, assuming a constant balance sheet χ , during normal periods, and solve for the price.

$$p_2(\lambda) = \begin{cases} 1, & \text{Issuer Solvent,} \\ \frac{1-\phi}{\lambda}, & \text{Issuer Insolvent,} \end{cases} \quad (5)$$

As transaction fees rise, the participation of arbitrageurs becomes limited, reducing the effective number of active arbitrageurs, n . Since $K \equiv \frac{1}{n\chi}$, a decline in n raises K . During stress periods, balance sheet capacity χ may also contract, further raising K . Together, these forces amplify price deviations and can prevent Terra's price from reverting to \$1.

I empirically test whether execution costs and fundamentals impede peg stabilization by regressing minute-level changes in the peg on measures of trading frictions and market controls. For the arbitrage mechanism to function effectively, price at the next timestamp must be pushed back to \$1. In my setting, I take open-to-open price change to measure arbitrage effectiveness. The dependent variable is the *open-to-open* price change, defined as $\Delta P_t = P_t^{\text{Open}} - P_{t-1}^{\text{Open}}$, and the regressors are constructed using a strictly lagged window $[t-6, t-1]$. Using open-to-open changes aligns the information set of the outcome with the predictors and mitigates concerns of simultaneity: P_t^{open} is observed at the start of minute t , before trades and fees realized *during* minute t can mechanically co-move with the left-hand side. Moreover, the open-to-open change captures the price jump that arbitrageurs inherit *before* they trade in minute t , providing a cleaner test of whether lagged variables predict subsequent peg restoration.

The baseline specification is

$$\begin{aligned}\Delta P_t = & \alpha + \beta_f \text{Fee Rate}_{[t-6,t-1]} + \beta_q \text{Tx Count}_{[t-6,t-1]} \\ & + \beta_L \ln\left(\frac{\text{LUNA}_{t-1}}{\text{LUNA}_{t-6}}\right) + \beta_B \ln\left(\frac{\text{BTC}_{t-1}}{\text{BTC}_{t-6}}\right) \\ & + \theta_{\text{day}(t)} + \theta_{\text{hour}(t)} + \varepsilon_t\end{aligned}\tag{6}$$

Here $\text{Fee Rate}_{[t-6,t-1]}$ is the transaction-weighted average fee, capturing the marginal cost of executing trades, and $\text{Tx Count}_{[t-6,t-1]}$ is the total number of transactions, capturing blockchain traffic, hence congestion. The LUNA return reflects ecosystem-specific risk, while the thirty-minute BTC return captures overall market sentiment in crypto, allowing news to be incorporated while remaining predetermined. Fixed effects $\theta_{\text{day}(t)}$ and $\theta_{\text{hour}(t)}$ (or calendar-hour on May 7) absorb time patterns; standard errors are clustered at the day or calendar-hour level.

To separate aggregate intensity from marginal cost, I also estimate

$$\begin{aligned}\Delta P_t = & \alpha + \beta_g \text{Gas Fees}_{[t-6,t-1]} + \beta_f^\perp \widetilde{\text{Fee Rate}}_{[t-6,t-1]} \\ & + \beta_L \ln\left(\frac{\text{LUNA}_{t-1}}{\text{LUNA}_{t-6}}\right) + \beta_B \ln\left(\frac{\text{BTC}_{t-1}}{\text{BTC}_{t-6}}\right) \\ & + \theta_{\text{day}(t)} + \theta_{\text{hour}(t)} + \varepsilon_t.\end{aligned}\tag{7}$$

where $\text{Gas Fees}_{[t-6,t-1]}$ is the total dollar fees paid on-chain, and $\widetilde{\text{Fee Rate}}_{[t-6,t-1]}$ is the residual from regressing fee rate on total fees, capturing variation in marginal cost orthogonal to size of activity. This allows us to test whether stabilization depends more on the scale of arbitrage trading or on the cost per transaction.

Formally,

$$\text{Gas Fees}_{[t-6,t-1]} = \sum_{\tau=t-6}^{t-1} \text{Gas Fees}_\tau, \quad \text{Fee Rate}_{[t-6,t-1]} = \frac{\sum_{\tau=t-6}^{t-1} \text{Gas Fees}_\tau}{\sum_{\tau=t-6}^{t-1} \text{TX}_\tau}.$$

The transaction count, TX_{τ} , includes all blockchain transfers, swaps, staking, and smart-contract calls. Retail traders account for most transactions by number, but the dollar amount of these trades is small compared to institutional traders, who transact less frequently but at much larger dollar values. If only a few institutional traders pay higher gas fees, the fee rate increases considerably. By contrast, when many retail transactions occur simultaneously, the gas fees increase across the board to secure timely processing of transactions. As a result, the average fee per transaction—the fee rate—increases, raising the marginal cost of trading.

All regressions are estimated at one-minute frequency across three non-overlapping windows: (i) May 1–6, 2022 (pre-run), (ii) May 7, 2022 (event day), and (iii) May 8–11, 2022 (post-run). Table 2 reports regressions of price changes on transaction costs, fundamental, LUNA’s price, and controls. Arbitrage fails when frictions overwhelm participation and fundamentals, making price convergence impossible. On May 7, two frictions interact to break down the arbitrage mechanism. While participation attempts to restore the peg—the Total Fees coefficient of $+2 \times 10^{-7}$ implies that an additional \$10,000 in fees predicts a +0.2 cents increase in ΔP_t , which translates to (+100,00 transaction counts +0.2 cent price correction). This positive pressure faces two offsetting frictions: Execution frictions through network congestion create immediate costs, the orthogonal fee rate coefficient $\widetilde{\text{Fee Rate}}_{\perp, [t-6, t-1]}$, of -0.00026 implies that a 0.01 increase in the fee rate every 5 minutes reduces ΔP_t by 0.00026 cents, such that an economically large increase of 7.7 units would be required to offset the participation benefit. Simultaneously, LUNA price declines create fundamental frictions: the coefficient of $+0.00365$ translates the day’s -12% LUNA return into a predicted -0.047 cents contribution to ΔP_t , meaning arbitrageurs faced adverse movements in the collateral asset during execution to correct the price of Terra. These combined frictions—marginal costs and collateral risk—overwhelmed participation efforts, causing persistent peg deviations and an effective breakdown of the arbitrage mechanism.

Post-May 7, the fundamental friction intensifies and dominates all other channels. Participation and execution-friction coefficients become economically small or statistically weak as arbitrageurs withdraw, while the LUNA return coefficient jumps to $+0.02429$ —a six-to-sevenfold increase. Now, 5-minute LUNA movements translate into median *tenths of cents*, and larger swings of ± 1 – 2 cents for UST. This represents fundamental friction at its extreme: the collateral asset's collapse makes arbitrage economically irrational regardless of execution costs, as traders would lock in losses even with perfect execution. Hence, participation vs. frictions is the key margin on May 7: activity pushed toward the peg, but cost frictions and weak fundamentals kept arbitrage from restoring the peg.

The progression from May 7's multi-friction binding (participation generating \$10,000 fees \rightarrow $+0.2$ cents, offset by marginal costs and LUNA price drop) to May 8–11's fundamental-friction dominance demonstrates that arbitrage requires all frictions to remain below critical thresholds. After May 7, fundamentals become the primary driver of UST price changes; fee-based channels are second-order. When execution frictions spike due to costs, fundamental frictions emerge from collateral volatility, or both bind simultaneously, arbitrageurs cannot profitably eliminate price discrepancies—the mechanism fails not from lack of activity to correct the prices but from the economic impossibility of profitable execution under binding constraints.

Table 2: Effect of Fees on Changes in Price Deviation

	ΔP_t		
	May 1–6	May 7	May 8–12
<i>Panel A: Fee Rate and Transaction Count</i>			
Fee Rate _[t-6,t-1]	-0.000001 (0.000003)	-0.00003 (0.00002)	0.00078 (0.00060)
Tx Count _[t-6,t-1]	-0.00000 (0.00000)	0.0000001*** (0.00000)	-0.0000000 (0.0000000)
BTC Return _[t-6,t-1]	0.00036 (0.00041)	-0.00717 (0.00737)	0.07023 (0.03326)
LUNA Return _[t-6,t-1]	-0.00001 (0.00022)	0.00292*** (0.00076)	0.02428*** (0.00353)
Observations	8,561	1,409	5,759
R ²	0.00021	0.05471	0.05430
<i>Panel B: Total Fees and Orthogonal Fee Rate</i>			
Total Fees _[t-6,t-1]	-0.00000 (0.00000)	0.0000002*** (0.0000000)	0.00000 (0.0000001)
Fee Rate _{⊥,[t-6,t-1]}	0.000001 (0.000004)	-0.00026*** (0.00007)	0.00081 (0.00065)
BTC Return _[t-6,t-1]	0.00036 (0.00043)	-0.00611 (0.00805)	0.07016 (0.03338)
LUNA Return _[t-6,t-1]	-0.00002 (0.00023)	0.00365*** (0.00052)	0.02429*** (0.00354)
Day FE	Yes	No	Yes
Hour FE	Yes	Yes	Yes
Observations	8,561	1,409	5,759
R ²	0.00020	0.05302	0.05423

Notes: This table reports regressions of peg-deviation changes on blockchain fees and controls over 5-minute intervals. Panel A uses Fee Rate_[t-6,t-1] and Tx Count_[t-6,t-1]. Panel B uses Total Fees_[t-6,t-1] and an orthogonalized Fee Rate, Fee Rate_{⊥,[t-6,t-1]}. On May 7, transaction volume surged, but rising marginal fees priced out participants and peg deviations persisted. Execution frictions outweighed participation, and the arbitrage mechanism failed to restore the peg. At the same time, LUNA's sharp price decline further hindered peg correction, reinforcing the breakdown of the arbitrage mechanism. After May 7, LUNA's collapsing price became the dominant driver of persistent peg deviations, reflecting failure of the arbitrage mechanism. All specifications include BTC and LUNA 5-minute returns. Fixed effects are at the day and hour levels; standard errors clustered accordingly.

5 Run Dynamics and Liquidity Collapse

One of the most widely used applications on the Terra platform was the Anchor protocol, which provided lenders with a stable interest of 19.5%. Anchor was by far the most important protocol in the UST network, accounting for 46% of the total network volume (Liu et al., 2023). Anchor functioned as a high-yield Terra’s bank, drawing many users because of the heavily subsidized deposit rates. The growth was stable, suggesting a strong investor belief in the sustainability of Anchor’s returns. Figure 6 displays the evolution of UST deposits in the Anchor Protocol from January through May 7, 2022. Deposits grew from under 2 billion USD and reached 14 billion on May 7. Given that Anchor accounted for over 70% of UST outstanding, understanding the dynamics within is critical. A high deposit base implies that sizable outflows from Anchor could put negative pressure on Terra’s stability, hence price. (Liu et al., 2023) show that the run on Terra on May 7th began with a few large investors withdrawing their UST deposits from Anchor and selling them on exchanges.

While (Liu et al., 2023) emphasize UST withdrawals by large investors, I study a complementary mechanism: the self-initiated withdrawal and protocol-triggered liquidations of collateral positions by borrowers. Figure 7 plots hourly burn and liquidation volumes of bLUNA and bETH—the two main bonded assets used to borrow UST on Anchor. A sharp rise in the burning and liquidations of bonded LUNA is visible a few hours after the arbitrage mechanism failed on May 7, indicating that borrowers rushed to redeem their collateral before the system’s stability deteriorated further. Investors anticipated falling LUNA prices and chose to unwind positions while they still could. At the same time, the surge in protocol-triggered liquidations reflects the deterioration of Anchor’s LTV ratios as falling collateral values mechanically pushed borrowers past threshold limits (80%). The

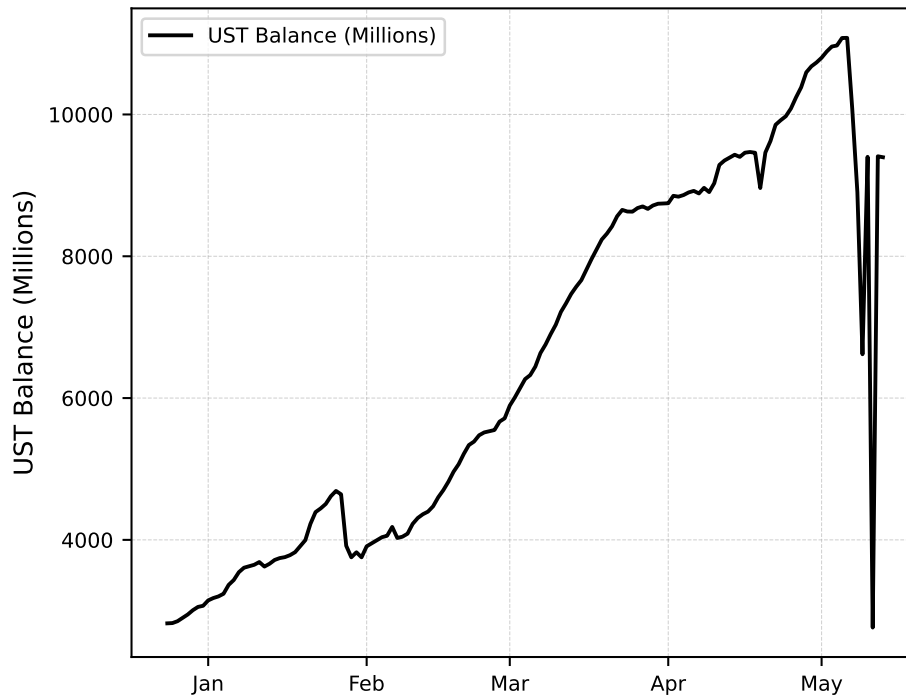


Figure 6: Daily Anchor UST balance – Token Count

burn and liquidation activity suggests that the run was initiated proactively, underscoring the fragility of the system after the arbitrage mechanism failed.

The preemptive burning reflects a belief that remaining in Anchor could lead to “inability to recover collateral” – for example, if LUNA’s price collapsed, any borrower who hadn’t exited might be liquidated at protocol-triggered fire-sales, potentially losing a large portion of their staked assets. Thus, burning activity became an early indicator of panic, capturing investors’ intent to flee the system while they still had the chance. By May 9, a significant wave of Anchor’s collateral had already been voluntarily pulled out. This underscores that the symptoms of the run started with voluntary exits: the sudden spike in bLUNA burns was essentially the initial stage of the run, driven by investors’ expectations of imminent collapse of LUNA and their desire to avoid forced liquidation losses.

In summary, the dramatic inflection soon after the arbitrage breakdown illustrates how the Anchor run was initiated by borrower behavior and the protocol's mechanism. The moment the stabilizing arbitrage failed and confidence wavered, bonded assets were converted back to LUNA en masse, shrinking Anchor's collateral base. The wave of user-initiated burning reflects a classic run dynamic: once the peg broke, early movers rushed to withdraw (burn) collateral, foreshadowing the cascade of liquidations and losses that would follow as slower actors were left behind. Such early, voluntary collateral withdrawals reveal the market's collective judgment that LUNA prices were no longer sustainable, marking the true start of the run.

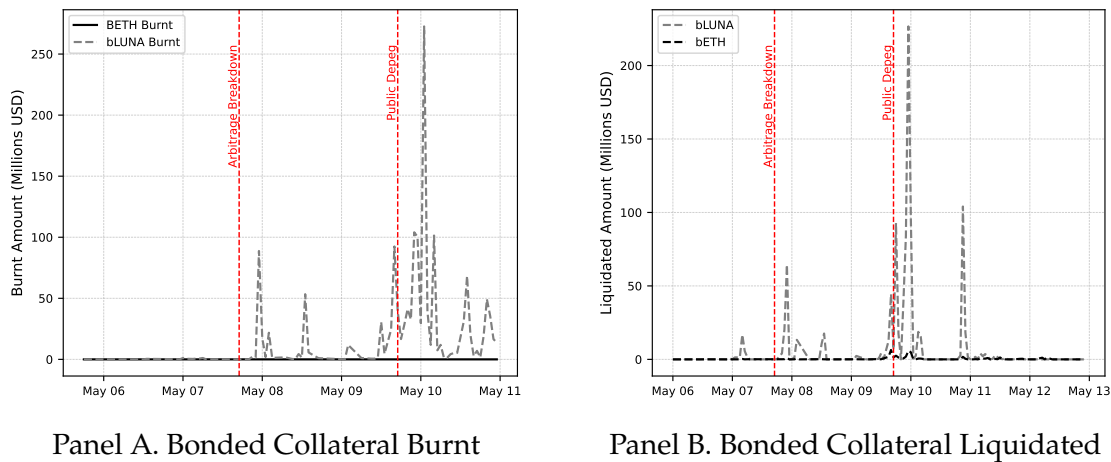


Figure 7: Burning and liquidation of bonded collateral assets (bLUNA and bETH) around the UST run. Vertical lines denote the arbitrage breakdown (May 7, 5 PM) and public depeg (May 9, 5 PM).

I further show the daily outflow magnitude of UST from Anchor. Figure 8 plots the net daily change in Anchor UST balances from May 1–9, 2022. The sharpest outflow occurred on May 7, when Anchor deposits fell by **1 billion UST**, over **7% reduction** in a single day. This initial withdrawal represented a direct liquidity shock to the broader ecosystem, as funds exited Anchor and flowed into both centralized and decentralized exchanges. Part of this outflow went to centralized

exchanges, where order books absorbed the sudden increase in sell pressure, while the remainder was routed to decentralized pools such as Curve. In the following days, the pace of withdrawals intensified: On May 8, an additional **1.2 billion UST** leave Anchor, followed by **2.3 billion UST** on May 9. These sequential outflows not only depleted Anchor's deposit base but also transmitted stress into secondary markets, where centralized exchanges' liquidity should be sufficient to absorb the huge selling pressure without a significant price impact.

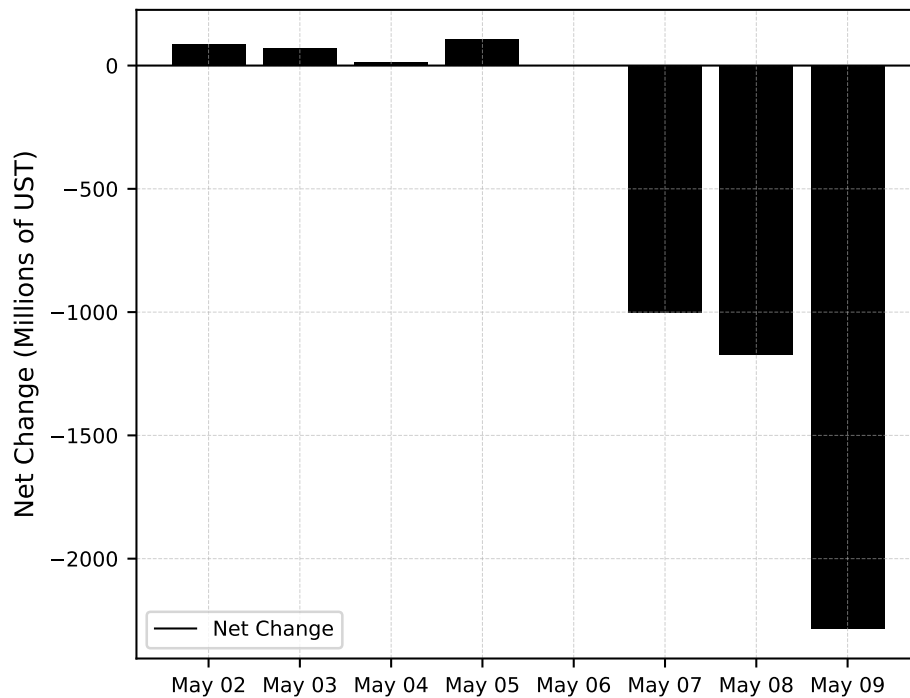


Figure 8: Daily Anchor UST balance Change

5.1 Price Impact and Liquidity in Centralized Exchanges

To assess how Anchor withdrawals transmitted stress to secondary markets, I analyze order book liquidity on the two primary venues for UST trading with fiat: Coinbase (UST/USD) and Binance (UST/USDT). Using minute-level order book snapshots, I simulate market orders of varying sizes (\$100,000 to \$5 million) to

measure the price impact of large withdrawals. This simulation approach captures execution feasibility more accurately than analyzing completed trades. During liquidity stress, large orders are typically fragmented into smaller parcels to minimize market impact, causing realized trade prices to understate true execution costs. By applying consistent order sizes (\$100 K- \$5 M) across all time periods, the volume-weighted average price (VWAP) simulations provide a standardized measure of market depth and price impact. This allows me to capture how much liquidity the UST / USDT or UST/USD order book could realistically supply during the run-up to the de-peg.

I calculate the VWAP from full depth-of-book snapshots. The methodology calculates VWAP by sequentially filling against available bid or ask depth until the target notional is reached. When visible depth proves insufficient to complete an order, the simulation returns missing values, indicating illiquidity at that size and time. I track how execution prices deviate from quoted mid-prices and identify periods when order books cannot absorb large transactions at any price. This analysis directly addresses the transmission mechanism from Anchor to secondary markets: as UST withdrawals flow from Anchor to exchanges, the order book analysis reveals whether centralized venue liquidity was sufficient to absorb the selling pressure without significant price dislocation. The results illuminate how macro-level funding shocks propagate through market microstructure.

5.2 Order Book Summary Statistics

Table 3 provides a summary statistics of order book depth on Binance and Coinbase. Both exchanges are widely used for trading activities, but Binance is more liquid than Coinbase. Binance operates as a global crypto exchange with a presence spanning over 180 countries, while Coinbase is primarily U.S.-based and is widely recognized as the largest cryptocurrency exchange in the United States.

During the normal period from May 1 to May 6, the Binance order book for the UST/USDT trading pair remained relatively balanced, with an average bid depth of approximately \$6.4 million and ask depth of \$6.6 million. The stress period (May 7) reveals pronounced asymmetry: bid depth contracted to \$3.7 million while ask depth expanded to \$10.7 million, consistent with intensified selling pressure and reduced market-making activity. The post-crisis period (May 8–9) exhibits partial normalization with ask depth declining to \$4.7 million and bid depth recovering to \$8.5 million, potentially reflecting liquidity provision interventions by Terraform Labs.

The Coinbase UST/USD market demonstrates substantially lower absolute liquidity levels. Pre-crisis conditions (May 1–6) show average bid and ask depths of \$409,000 and \$427,000, respectively, with median bid depth at \$388,000. During the stress period, while average bid depth remained relatively stable at \$368,000, ask-side liquidity experienced significant deterioration: median ask depth fell to \$266,000 with minimum ask levels reaching \$1,200. Post-crisis conditions reveal further fragmentation with average depths declining to \$273,000 (bid) and \$177,000 (ask), and both minimum bid and ask levels falling below \$500.

These patterns document a systematic migration from deeper to shallower liquidity pools as market stress intensifies, consistent with theoretical predictions of liquidity provision under adverse selection and inventory risk ([Glosten and Milgrom, 1985](#)). Under heightened information asymmetry, market makers widen spreads and reduce quoted depths to mitigate adverse selection costs, while inventory-constrained dealers curtail position-taking capacity as volatility increases.

Table 3: Top-20 Depth for Binance (UST/USDT) and Coinbase (UST/USD)

Exchange	Period	Mean	Median	Min	Max
Binance Bid	May 1–6	\$6.43M	\$6.35M	\$1.47M	\$13.34M
	May 7	\$3.72M	\$3.33M	\$144K	\$10.63M
	May 8–9	\$8.46M	\$7.70M	\$26.36K	\$30.55M
Binance Ask	May 1–6	\$6.60M	\$6.32M	\$2.27M	\$16.66M
	May 7	\$10.67M	\$11.47M	\$130K	\$19.51M
	May 8–9	\$4.68M	\$4.78M	\$16.25K	\$16.55M
Coinbase Bid	May 1–6	\$409K	\$388K	\$4.27K	\$2.09M
	May 7	\$368K	\$388K	\$5.00K	\$793K
	May 8–9	\$273K	\$204K	\$438	\$974K
Coinbase Ask	May 1–6	\$427K	\$379K	\$9.51K	\$1.59M
	May 7	\$386K	\$266K	\$1.23K	\$1.16M
	May 8–9	\$177K	\$122K	\$138	\$1.58M

5.3 Liquidity Dynamics on Binance

5.3.1 Institutional Order Execution Analysis (\$2–\$10 Million)

I simulate market order execution for institutional-scale transactions up to \$10 million on Binance’s UST/USDT pair, representing approximately one percent of the initial \$1 billion Anchor withdrawal. Figure 9 presents volume-weighted average prices (VWAPs) for three hypothetical institutional orders—\$2 m (blue), \$5 m (orange), and \$10 m (green)—from 00:00 UTC on 7 May through 09:00 UTC on 9 May, distinguishing between sell-side sweeps (bid execution) and buy-side sweeps (ask execution).

Pre-stress liquidity asymmetries (00:00–17:40 UTC, May 7.) Buy-side execution demonstrates substantial depth across all order sizes, with execution prices within approximately 2 basis points of parity. Conversely, sell-side liquidity exhibits severe constraints: \$10 million orders consistently fail to execute, returning null values throughout this period. This asymmetry indicates that bid-side depth

constraints preceded visible market stress, suggesting structural vulnerability to large redemption flows. The \$2 million orders execute near parity, while \$5 million orders incur 3–4 basis points of price impact even under normal conditions.

Liquidity deterioration following Anchor outflows. The initial large on-chain UST deposit (\$175 million at 17:45 UTC) coincides with immediate structural breakdown in order book depth. Bid-side capacity for \$5 million orders disappears instantaneously, restricting executable block size to \$2 million. Thirty eight minutes later, following the subsequent \$35 million withdrawal, even \$2 million orders exhibit intermittent execution failures, with successful fills occurring at prices as low as \$0.9855 (145 basis points). At this point, the Binance order book became incapable of absorbing any institutional-scale UST redemption flow, effectively severing the critical link for peg defense.

Run Dynamics on Binance. The systematic disappearance of liquidity tiers provides evidence of run dynamics consistent with theoretical predictions. Rather than gradual depth deterioration through progressive price concessions, I observe abrupt capacity truncation above a declining ceiling that responds directly to cumulative redemption volumes. The \$10 million tier remains consistently un-executable, followed by \$5 million tier elimination, and finally \$2 million tier intermittent failure following cumulative on-chain redemptions exceeding \$20 million. The fact that even a single \$10 m redemption could not be sold at any point during the window quantifies the magnitude of the imbalance between the sellers willing to sell UST and the buyers willing to take sell orders. This microstructural evidence documents the precise transmission mechanism through which Anchor withdrawals propagated to secondary market. The results, therefore, tighten the evidence for classic run dynamic on the Binance. The peg failure is suggestive of the price impact of Anchor inflows to CEX and the absence of buyers to take sell orders. Figure 9 shows the liquidity and price impact dynamics for large orders in Binance.

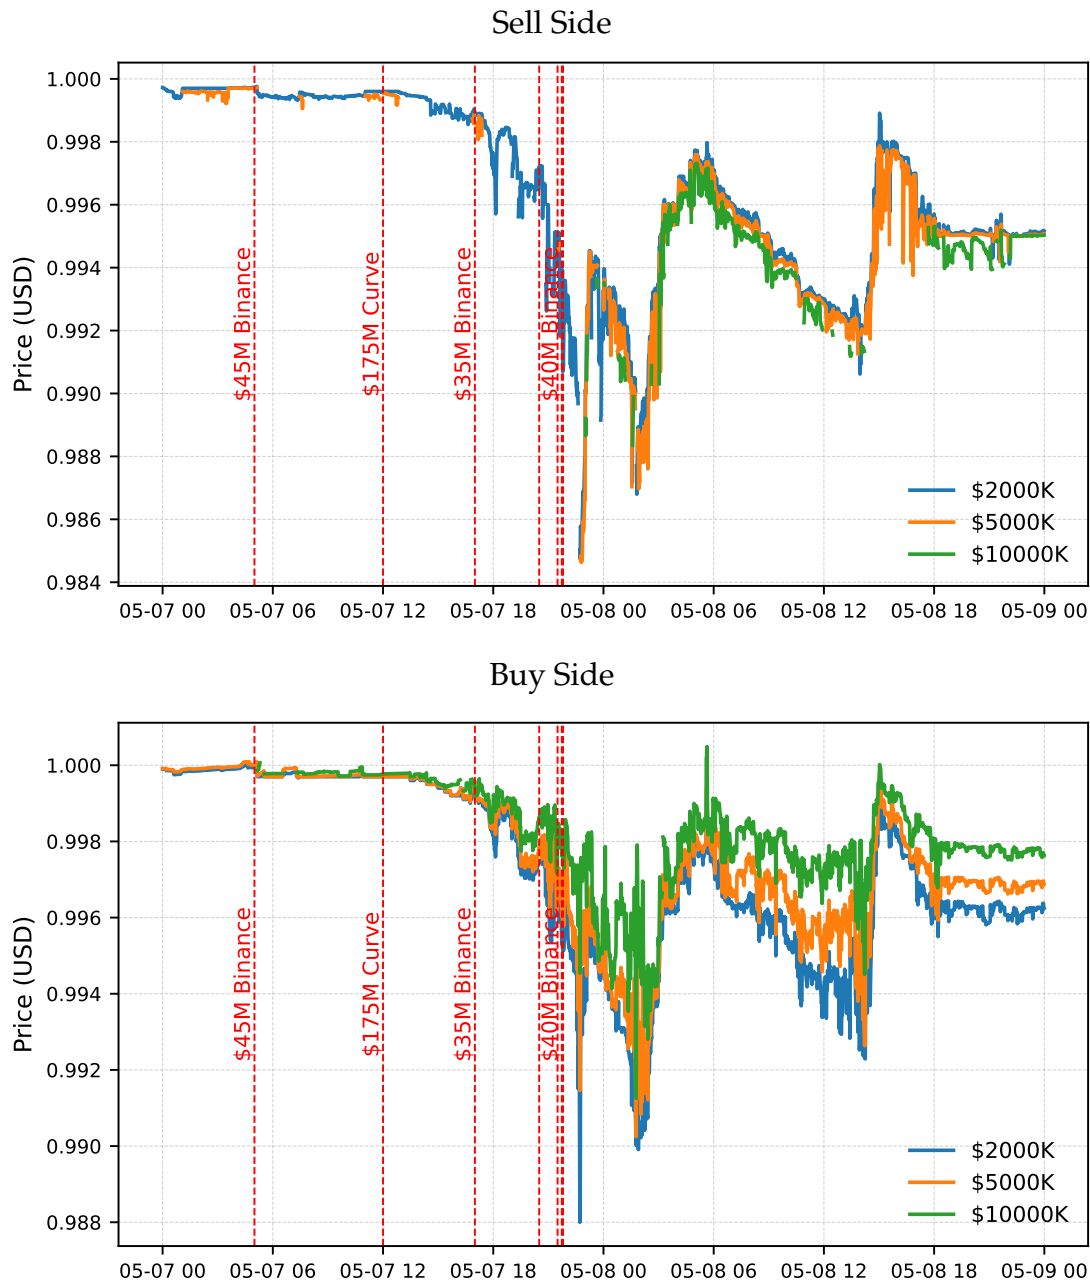


Figure 9: Binance *UST/USDT* limit-order book, 7–9 May 2022. Dashed vertical lines label cumulative Anchor withdrawals to various destinations. Missing segments indicate minutes in which the book cannot absorb the trades. The sell-side fractured plot indicates that buy-side liquidity was absent in Binance. Buy Side plot indicates that sell - side liquidity was available in Binance. This plot represents the microstructure of a run. Trades from \$2 - \$10 million were swept through the order book.

5.3.2 Run Dynamics on Coinbase

I evaluate execution quality for mid-size institutional orders on Coinbase's UST/USD pair, focusing on \$500 k and \$1 million market orders during the May 7–8 liquidity crisis. Unlike Binance, Coinbase's shallower order book could not absorb orders exceeding \$1 million, necessitating analysis of smaller institutional flows to understand venue-specific liquidity dynamics.

Early Stability with Size-Dependent Fragility (May 7, 00:00–18:00 UTC). The first panel reveals an important asymmetry in Coinbase's liquidity provision during the initial stable phase. While \$100 k orders execute essentially at par throughout most of May 7, \$1 million trades exhibit visible slippage by 18:00 UTC—several hours before similar deterioration appears in smaller order sizes. This kinked response pattern is consistent with a depth curve that flattens first at the tail, effectively rationing the largest institutional flows while preserving execution quality for retail-scale transactions.

Symmetric Order Book Deterioration (May 7, 18:00–24:00 UTC). The buy-side panel demonstrates that ask-side depth remained ample through most of May 7, with \$500 k–\$1 m orders executing within a few basis points of parity. However, after approximately 18:00 UTC, all VWAP curves drop in near-parallel fashion: \$1 million buy orders print as low as \$0.990, while even \$500 k purchases clear at noticeable discounts. The step-wise, size-invariant deterioration indicates that multiple liquidity layers were withdrawn simultaneously rather than spreads merely widening at the top of book.

Bid-Side Collapse and Order Book Hollowing (May 7, Evening–May 8). The sell-side panel reveals a mirror image of buy-side stress. Early in the day, \$1 million market sales achieve only modest discounts, but as Anchor redemptions accumulate, the bid stack fractures dramatically. By evening May 7, identical orders

clear below \$0.986—representing over 140 basis points of price impact—while even \$500 k blocks suffer double-digit basis point slippage.

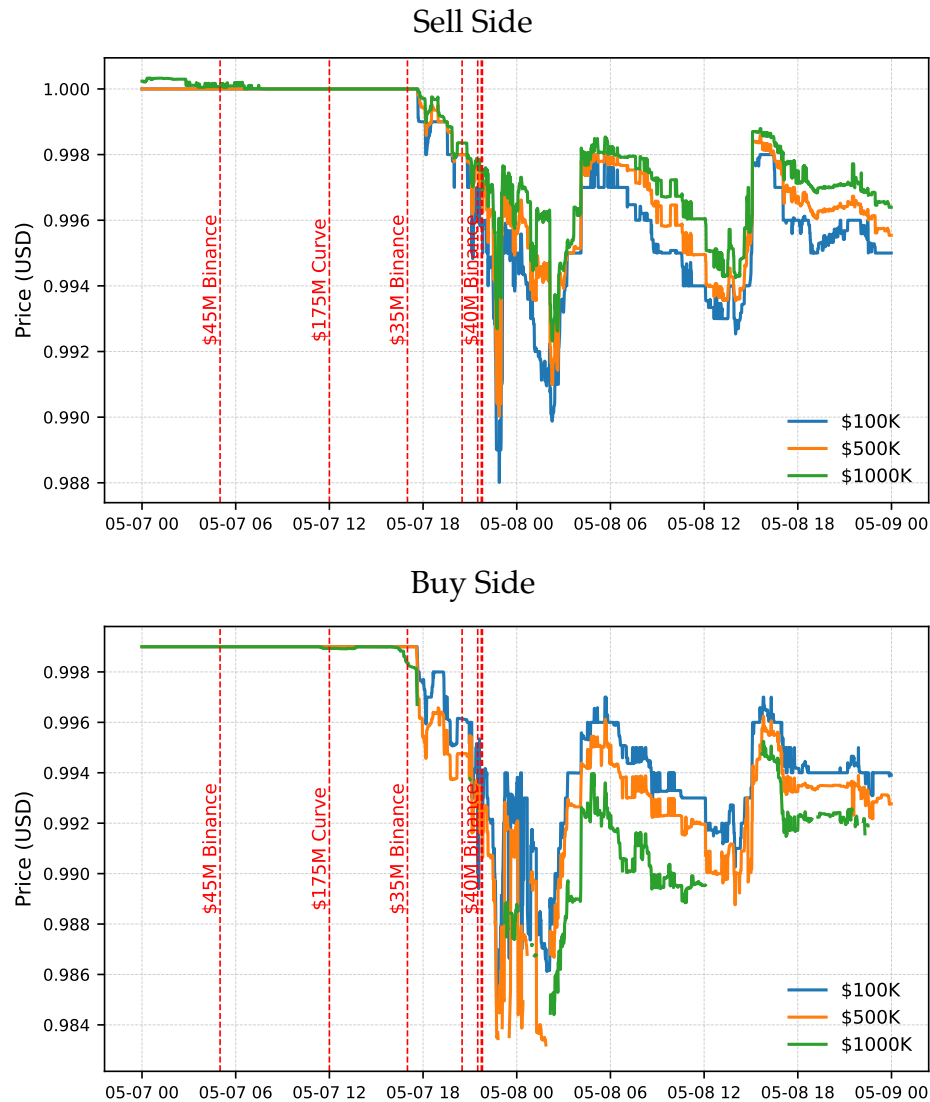


Figure 10: VWAP execution prices for simulated market orders of \$100k–\$1m in the Coinbase *UST/USD* limit-order book, 7–9 May 2022. Dashed vertical lines mark large Anchor withdrawals routed to centralized and decentralized venues. Missing segments indicate minutes in which the visible book could not supply the specified notional at any price. On the sell side, the absence of quotes shows that buy-side liquidity had vanished. On the buy side, quotes persisted, indicating that sellers on Coinbase could still find counterparties. Together, the panels depict the microstructure dynamics of the run, as successive market orders swept available depth and destabilized prices.

5.3.3 Cross-Venue Liquidity Drain.

Figure 11 places these developments in a two-day perspective by tracking the \$1 m buy- and sell-side VWAPs alongside the quoted mid-price from May 6 to May 8. Until approximately 17:00 UTC on May 7, both VWAPs tracked close to the observed price, implying that \$1 m trades could still be executed near par. After 17:00 UTC, sell side liquidity deteriorated sharply on Coinbase: the sell-side VWAP (red line) plunged more than 100 bp before disappearing entirely, while Binance maintained continuous liquidity throughout. Critically, the missing sell-side observations on Coinbase after 20:00 UTC reveal that visible depth was insufficient to fill a \$1 m order at any price, while Binance continued to provide liquidity for institutional-sized trades. The figures demonstrate that Coinbase's order book failed completely for \$500k–\$1m trades within hours of arbitrage failure on May 7, well before the peg collapsed on May 9, while Binance remained functional as the market's liquidity backstop.

Chronology of failure. On Coinbase, the \$1m sell-side VWAP breaks the 100-bp barrier by 17:45 UTC and plummets to \$0.985 by 20:00 UTC before liquidity vanishes entirely—evidenced by the missing red line. Binance on the other hand maintained continuous liquidity for \$1 million trades throughout this period, with sell-side VWAP deteriorating more gradually to \$0.987. Thus, Binance acted as the market's last line of defense, continuing to absorb institutional redemption flows even as Coinbase's order book completely failed.

Cross-venue liquidity. The differential resilience is striking: Coinbase's complete inability to fill \$1 million orders (missing data points) versus Binance's continued market-making (continuous lines) demonstrates how liquidity concentrated in the deepest venue during the run. This pattern reveals that smaller exchanges face binding liquidity constraints during redemption cascades, with

their order books evaporating entirely while larger venues continue functioning, albeit with degraded execution quality.




Figure 11: Price Impact and Liquidity: Coinbase vs. Binance. This figure overlays VWAP execution prices for \$1 million market orders on the *UST/USD* order book of Coinbase and the *UST/USDT* order book of Binance, from 6–9 May 2022. Dashed red vertical lines denote cumulative Anchor withdrawals to various venues. Gaps in either the buy- or sell-side series indicate minutes when the order book lacked sufficient depth to absorb the quoted trade size. The sell side fractured on Coinbase—indicating absent buy-side liquidity—while Binance retained sell-side support. Market orders of \$1 million swept through the visible book, revealing where liquidity first vanished.

5.4 Liquidity Drain on Smaller Exchanges

A natural concern is that the liquidity contraction observed on Coinbase and Binance might be venue-specific or depend on the ticket sizes chosen. To address this, I replicate the trade simulation on two additional centralized exchanges with distinct clientele and market structures:

- **Huobi**, an Asia-centered platform where the UST/USDT book is retail-driven and never exceeds \$500k depth.
- **Kraken**, a U.S.-regulated exchange with historically thinner stablecoin liquidity but fiat settlement.

All procedures and redemption markers follow those used for Binance. This covers four major CEXs across two continents and removes selection concerns. Both Huobi and Kraken also exhibit depth collapse: neither offers economically viable execution, ruling out the possibility that buyers could shift elsewhere. These findings suggest the run was *global*. Figures are reported in Appendix B.

5.4.1 Huobi: Depth Vanishes Below \$500k

Even at a reduced notional of \$500k, Figure 2B shows sell-side VWAP falling from 0.997 to 0.988 within two hours after the \$175m Anchor withdrawal to Curve, with several minute-wide gaps of zero visible liquidity. The buy side deteriorates in parallel, briefly spiking above par before collapsing, and spreads widen to nearly 90 bp. Huobi therefore offers no viable buyers even for modest sell-side trades.

5.4.2 Kraken: Chronic Thinness, Rapid Collapse

On Kraken, a \$1m buy already required a 3–4 bp concession on May 6, while the sell-side VWAP was 35 bp below market. After redemptions exceeded \$35m

(22:00 UTC, May 7), sell-side execution fell to 0.986 and soon became unfilled. Single-sided panels (Figure 3B) show a stepwise withdrawal of buyers while sellers remained, confirming the run.

5.5 Cross-Venue Synthesis

Table 4: Capacity thresholds and collapse timing across exchanges

Venue	Max size	Spread	Buy VWAP	Sell VWAP
Coinbase	\$1m	<2 bp	0.993	0.986
Binance	\$5m	<2 bp	0.994	0.984
Huobi	\$0.5m	~4 bp	0.995	0.988
Kraken	\$1m	~40 bp	0.996	0.986

Table 4 highlights three stylized facts:

1. **Monotone ceilings:** failure thresholds align with the size of Exchange—Huobi at \$0.5m, Kraken at \$1m, Coinbase above \$1m, and Binance at \$2–5m.
2. **Synchronized timing:** all venues lose executable depth within a ± 30 -minute window of the second redemption spike, ruling out idiosyncratic outages.
3. **Global quantity gate:** every book records minute-long intervals of *illiquidity*, indicating a system-wide run.

In sum, replication on Huobi and Kraken confirms the run was *global* rather than venue-specific. The synchronized timing and simultaneous depth collapse across Asia-centered and U.S.-regulated exchanges suggest a highly interconnected and possibly algorithmic market-making response. Liquidity providers, regardless of clientele or geography, appear to share similar risk management strategies or react to common information (on-chain redemptions, LUNA price collapse), leading to a collective withdrawal of buyers. This implies systemic risk: diversification across exchanges offers limited protection against ecosystem-wide liquidity shocks.

5.6 System-Wide Liquidity Stress Test

To assess the aggregate market capacity for UST trading during the crisis, I construct a comprehensive measure of system-wide executable liquidity. The analysis merges order-book snapshots for every UST-denominated pair (UST/USDT, UST/USD, UST/BTC, and others) traded on ten major centralized exchanges—Binance, Bittrex, Bibox, Coinbase, Gemini, Kraken, Huobi, KuCoin, OKX, Poloniex, and FTX—covering the four-day window from May 5, 00:00 UTC to May 8, 23:59 UTC. Data timestamped in milliseconds is converted to UTC datetimes and bucketed to the nearest minute, retaining only the most recent quote update within each minute for every venue. This creates a synthetic order book representing an optimistic upper bound on executable depth, as it assumes instantaneous routing across venues with zero frictions. I then simulate market orders of \$20 million, \$30 million, and \$35 million on both buy and sell sides to measure price impact following the algorithm detailed in Appendix 7.

Table 5 reveals the extreme illiquidity of UST markets during the May 5–8, 2022 crisis window. The results are stark: out of 11,520 one-minute snapshots, a \$20 million market sell order could be fully executed in only 28 instances—representing 0.24% of the observation period. During these rare moments when \$20 million orders were executable, the average VWAP was \$0.9974, with substantial variation (standard deviation of 0.0017). Critically, the minimum observed execution price fell to \$0.9949, indicating that even when liquidity existed, large orders faced price impacts exceeding 50 basis points from par.

For larger order sizes, the market structure constraints become absolute: both \$30 million and \$35 million orders show zero executable instances across the entire four-day period. The complete absence of observations indicates that the aggregate order book across all ten exchanges never contained sufficient depth to absorb a single \$30 million trade. This finding establishes that the market’s

effective absorption capacity was capped below \$30 million—orders of magnitude below the \$1 billion daily Anchor outflows. With redemption demand exceeding available liquidity by a factor of 30-50, the absence of buyers on centralized exchanges prevented the arbitrage mechanism from functioning: without buy-side depth to absorb selling pressure that created price impact, arbitrageurs could not purchase discounted UST on exchanges and route it back to primary exchange for guaranteed \$1.00 redemption. This liquidity constraint in centralized exchanges, major liquidity providers, acted as friction preventing peg correction. Results provide microstructure evidence of a run and that the run on centralized exchanges caused UST’s further collapse.

Table 5: Summary Statistics of VWAP by Order Size (May 5–8, 2022)

Statistic	VWAP \$20 M	VWAP \$30 M	VWAP \$ 35 M
Count	28	0	0
Mean	0.997445	–	–
Std Dev	0.001655	–	–
Min	0.994919	–	–
25%	0.995931	–	–
Median	0.997937	–	–
75%	0.998968	–	–
Max	0.999471	–	–

Notes: VWAP computed from aggregated order books across ten major exchanges. Missing values indicate insufficient depth to execute orders of specified size across all venues combined.

6 Robustness Test

On May 12, 2022, Tether—the world’s largest stablecoin—briefly depegged to \$0.95 during peak UST market stress before regaining parity within 24 hours⁸. Over \$4

⁸<https://www.reuters.com/markets/us/crypto-collapse-intensifies-stablecoin-tether-slides-below-dollar-peg-2022-05-12/>

billion flowed from UST into Tether (Figure 5B), yet Tether did not collapse. This episode provides a natural benchmark for order-book resilience.

I restrict the analysis to Coinbase, which offers detailed order-book data for the USDT–USDC pair. Figure 4B shows that at the height of the Tether depeg, a \$2m sell order could not be executed, implying visible depth below \$2m. By 23:10 UTC on May 14, the same order cleared at \$0.9865 (−1.35% slippage), and within 24 hours slippage tightened to below 0.3%.

By contrast, an identical \$2m sell of UST into USDT on Binance on May 8 generated a 9.8% price impact and remained unfilled for nearly two days. The brief and shallow illiquidity for USDT versus the prolonged collapse for UST corroborates the central mechanism: fully collateralized stablecoins exhibit only transient depth shrinkage, while stablecoin designs dependent on weak fundamentals, such as Terra’s, sustain exhaustion of order-book capital resulting in full-blown run.

7 Conclusion

This paper documents the breakdown of Terra’s arbitrage mechanism two days before its collapse. Despite profitable opportunities, market frictions and weakening fundamentals deterred participation, preventing peg correction. Once arbitrage failed, liquidity evaporated on centralized exchanges, buyers disappeared, and a classic run dynamic unfolded, driving the collapse of stablecoin, Terra.

The findings provide the first microstructure evidence consistent with theoretical models of the arbitrage–run tradeoff: broad participation stabilizes prices in normal times but amplifies run probability (Ma et al., 2023). Importantly, execution frictions—often considered second-order—can become first-order drivers of systemic instability when they hinder arbitrage. Monitoring such frictions offers

early-warning indicators with implications for financial stability surveillance, market design, and the regulation of arbitrage-dependent instruments.

The framework developed here generalizes beyond stablecoins. Arbitrage breakdowns also characterize money market funds “breaking the buck” or ETF dislocations (e.g., bond ETFs in March 2020). Applying this methodology to other fixed-value assets can shed light on common failure modes across markets.

Finally, the results carry direct policy implications. While the GENIUS Act addresses reserve soundness, stability also requires sufficient executable liquidity on exchanges. Just as central banks defending a currency peg must ensure reserve capacity, decentralized systems may need mechanisms to guarantee depth at scale, potentially through regulatory thresholds or mandated interventions. Can the proposed interventions prevent arbitrage failure in safe assets and the run that follows? I expect ample opportunities for future work to examine the implications of the policy interventions.

References

- Anadu, K., Azar, P., Huang, C., Cipriani, M., Eisenbach, T. M., La Spada, G., Landoni, M., Macchiavelli, M., Malfroy-Camine, A., and Wang, J. C. (2023). Runs and flights to safety: Are stablecoins the new money market funds?
- Capponi, A., Jia, R., and Yu, S. (2024). Price discovery on decentralized exchanges. *Available at SSRN 4236993*.
- Catalini, C. and de Gortari, A. (2021). On the economic design of stablecoins. *Available at SSRN 3899499*.

- d'Avernas, A., Maurin, V., and Vandeweyer, Q. (2022). Can stablecoins be stable? *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2022-131).
- Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics*, 14(1):71–100.
- Goldstein, I. and Pauzner, A. (2005). Demand–deposit contracts and the probability of bank runs. *the Journal of Finance*, 60(3):1293–1327.
- Gorton, G. B. and Zhang, J. Y. (2023). Taming wildcat stablecoins. *U. Chi. L. Rev.*, 90:909.
- Gromb, D. and Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of financial Economics*, 66(2-3):361–407.
- 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.
- Ma, Y., Zeng, Y., and Zhang, A. L. (2023). Stablecoin runs and the centralization of arbitrage. *Available at SSRN* 4398546.
- Makarov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319.
- Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *The Journal of finance*, 52(1):35–55.

Uhlig, H. (2022). A luna-tic stablecoin crash. Technical report, National Bureau of Economic Research.

Appendix A. VWAP Price Impact Algorithm

Algorithm 1 VWAP-Based Price Impact for Order Book Snapshots

Require: Order book snapshots \mathcal{B}_t for minutes $t \in T$

Require: Order sizes $S \in \{S_1, S_2, \dots, S_k\}$

Ensure: $\text{VWAP}_{t,\sigma}^{(S)}$ for each t , side σ , and size S

```

1: for each minute  $t \in T$  do
2:   Extract snapshot  $\mathcal{B}_t$  at latest timestamp in minute  $t$ 
3:   for each side  $\sigma \in \{\text{buy}, \text{sell}\}$  do
4:     Let  $\mathcal{A}_t \leftarrow$  asks if  $\sigma = \text{buy}$ , bids otherwise
5:     Sort  $\mathcal{A}_t$  by ascending price if  $\sigma = \text{buy}$ , descending otherwise
6:     for each order size  $S \in \{S_1, \dots, S_k\}$  do
7:        $R \leftarrow S, Q \leftarrow 0, V \leftarrow 0$ 
8:       for each level  $(p_i, q_i) \in \mathcal{A}_t$  do
9:          $n_i \leftarrow \min(R, p_i \cdot q_i)$ 
10:         $Q \leftarrow Q + \frac{n_i}{p_i}$ 
11:         $V \leftarrow V + n_i$ 
12:         $R \leftarrow R - n_i$ 
13:        if  $R \leq 0$  then
14:          break
15:        end if
16:      end for
17:      if  $R = 0$  then
18:         $\text{VWAP}_{t,\sigma}^{(S)} \leftarrow \frac{V}{Q}$ 
19:      else
20:         $\text{VWAP}_{t,\sigma}^{(S)} \leftarrow \text{NaN}$ 
21:      end if
22:    end for
23:  end for
24: end for

```

Appendix B. Figures

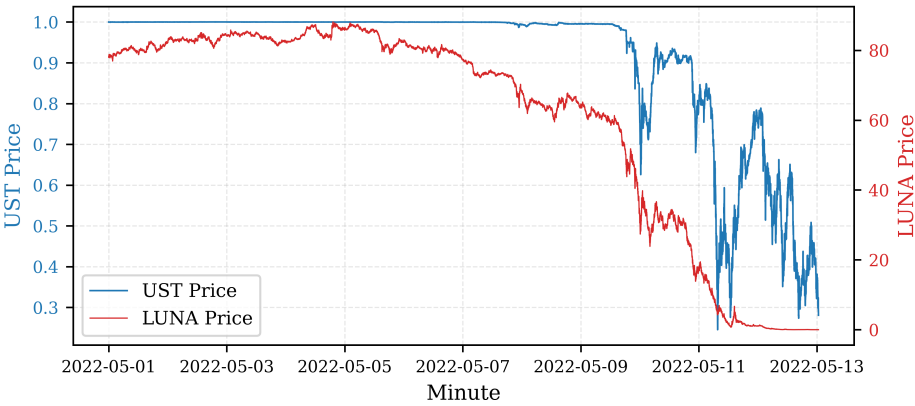


Figure 1B: UST-LUNA Price.

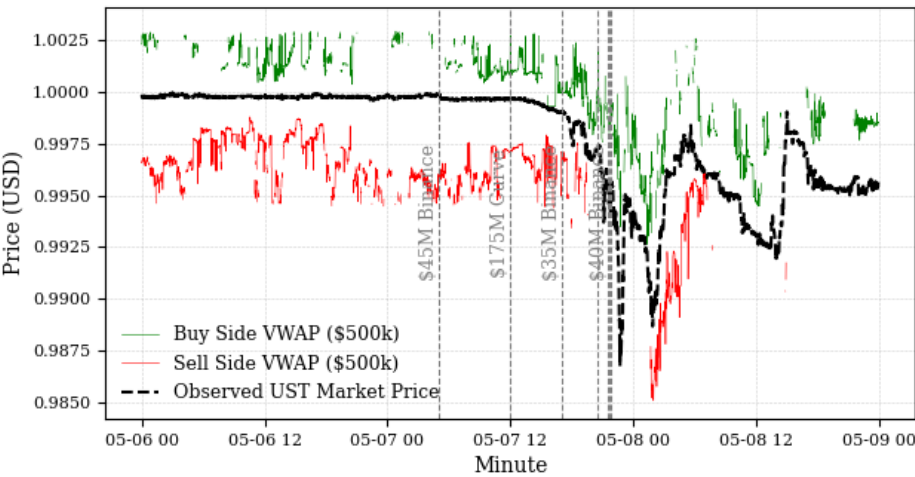


Figure 2B: Huboi Liquidity Dynamics and Price Impact.

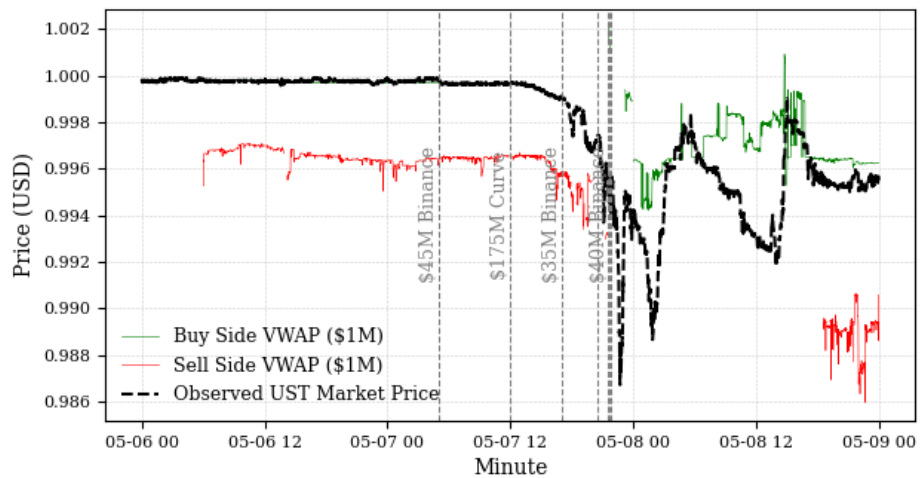


Figure 3B: Kraken Liquidity Dynamics and Price Impact.

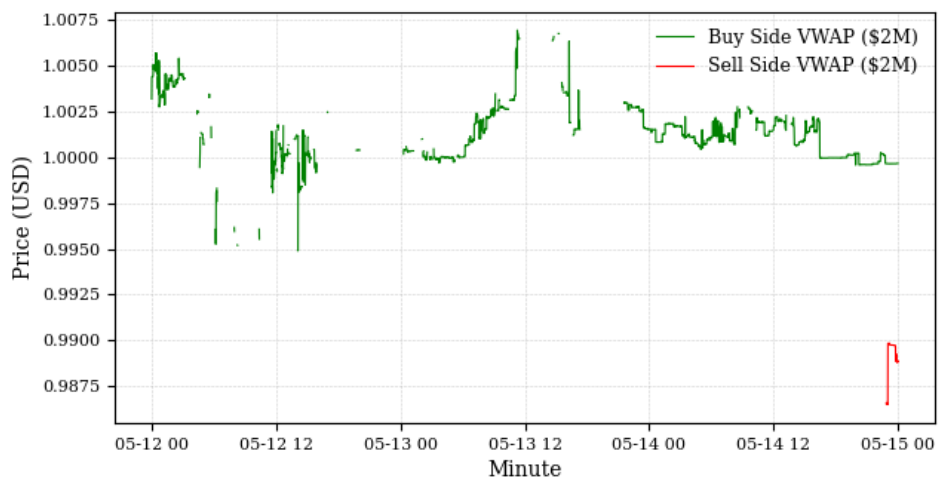


Figure 4B: Tether (USDT) Liquidity Dynamics and Price Impact.

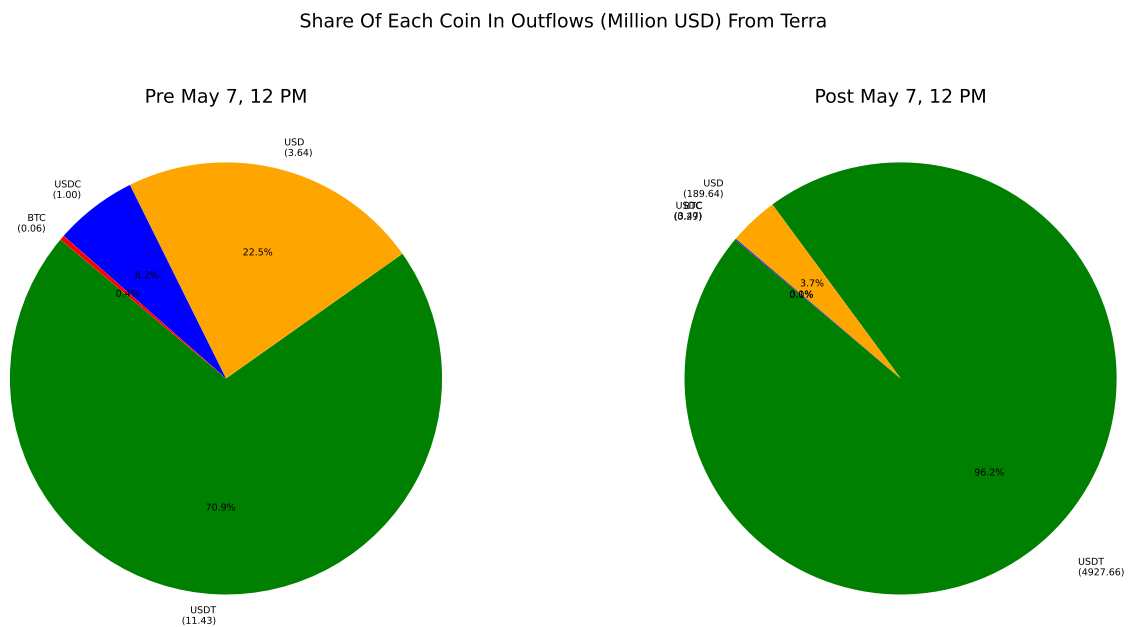


Figure 5B: Net dollar outflows from Terra around collapse: This figure shows the net outflows from Terra during the period: May 1, 2022, to May 13, 2022. USDT shows the highest dollar inflows of \$ 4 billion from Terra.