Limits to Arbitrage and Runs on Stablecoins

Ravi Joshi*

August 28, 2024

Abstract

Limits to arbitrage may cause prices to deviate from fundamentals. In stablecoins, which rely on arbitrage to maintain their peg, they can be a source of run risk. Using tick-by-tick data on Terra, I show that arbitrage failure resulted in Terra depegging from $1 on May 7, 2022, following which the run began on May 9. A self-fulfilling panic began where negative sentiment and disagreement among investors increased, adverse selection risk rose, and eventually, Terra became almost worthless by May 12. These dynamics highlight limits to arbitrage as an additional source of run risk for safe assets.

Keywords: Limits to Arbitrage, Adverse Selection, Run, Flight to Safety. (JEL Codes: G12, D82, G21, G01, G23)

*Department of Finance, Louisiana State University, rjosh15@lsu.edu

1 Introduction

Stablecoins are cryptocurrencies designed to maintain a stable value of $1 through an arbitrage mechanism: when the price of a stablecoin deviates from $1, arbitrageurs can buy it at the market price and redeem its nominal value from its issuer or vice versa. They were introduced to serve as reliable stores of value and mediums of exchange in the volatile world of cryptocurrencies. However, like bank deposits, stablecoins are susceptible to runs. In May 2022, a stablecoin named *Terra* experienced a run which led to its eventual collapse. I show that a novel mechanism, common among all stablecoins and other traded safe assets, was behind Terra's crash.

(Diamond and Dybvig, 1983) introduced the canonical model of runs where depositors rush to withdraw their deposits because they fear that others might withdraw first. The bank has to sell its assets at a loss to accommodate the withdrawals, which actually renders the bank insolvent. Goldstein and Pauzner (2005) further illustrated that information can precipitate runs. In their framework, depositors observe a noisy signal about the asset's fundamentals, whose low realization can spark a run. When an asset is publicly traded, as stablecoins are, its price is always observable and might temporarily deviate from the nominal value. In the presence of limits to arbitrage, these deviations might persist longer than expected and act as the signal in Goldstein and Pauzner (2005). Investors do not know the source of these deviations: they might be due to weak fundamentals or temporary inefficiencies in the market. Hence, observing price deviations increases the probability of weak fundamentals, which leads to a standard run.

I demonstrate this theory by closely examining the sequence of events leading to Terra's crash. In a panic-based run, we expect several developments before the arbitrage mechanism breaks down. First, there would be significant capital withdrawals as investors flee to other assets. Second, adverse selection would increase as some investors were more informed about the asset value. Third, negative information would spread, and disagreement would rise until the investors reached a consensus about the asset's instability. However, my empirical findings show that these events —capital outflows, increased adverse selection, and the spread of negative information— all occurred *after* the breakdown of the arbitrage mechanism.

I begin my analysis by identifying when Terra's arbitrage mechanism broke down. When the arbitrage mechanism works well, arbitrage profits, measured as the absolute price deviations from \$1, shrink toward zero over time. Formally, the AR(1) coefficient for arbitrage profits should be less than one. Using trading data from Binance, I test this hypothesis for every hour of trading in May 2022. The AR(1) coefficient remained below one until May 7, when, following a large withdrawal (\$175M), it rose to one and stayed close to it throughout the crisis period. As a result, Terra's price deviated from its par value of \$1 around the same time by a small amount (0.367 bps) difficult to distinguish from its regular price deviations.

If a panic had initiated the run on Terra, we would have observed large outflows before the arbitrage breakdown. I calculate Terra's net flows for every hour of trading in May 2022. Contrary to this hypothesis, net flows were close to zero during the first week of May 2022. From May 7 to May 9, the net outflow to Terra was positive and became negative after May 10. Thus, I find no evidence for panic-based withdrawals before the arbitrage mechanism broke down on May 7. After the mechanism breaks down, we document major net flows from Terra. We observe a \$23.4 million outflow from *Terra* on May 8 at 3 AM, followed by a \$22 million withdrawal at 4 PM. On May 9, a significant withdrawal of \$118.58 million occurs at 5 PM, followed by \$38 million at 10 PM and \$46.34 million at 11 PM. On May 10 at 12 AM, we first observe a significant inflow to *Terra* of \$62.7 million, followed by several additional inflows. The highest inflow is recorded on May 11 at 7 AM, amounting to $461.7 million, followed by several other inflows in a short time. However, the trend starts to reverse, and we observe outflows after May 11.

My second hypothesis indicates that if the run on Terra were induced by a panic, adverse selection would have risen before the arbitrage breakdown. To test this hypothesis, I follow Glosten and Harris (1988) to estimate Terra's adverse selection, measured as a component of the bid-ask spread that compensates the market maker for taking on the adverse selection risk per one dollar of the traded asset. Stablecoins have much lower adverse selection compared to other cryptocurrencies, which is consistent with their role as safe assets. Second, I find that Terra's adverse selection rose during the crisis and stayed elevated afterward. In contrast, adverse selection did not rise significantly for other cryptocurrencies, including stablecoins, indicating little to no spillover from Terra. Most importantly, adverse selection only rose after the large withdrawal on May 7 and not before, rejecting a panic-based run.

Finally, I explore the role of information diffusion, particularly through social media, in the Terra-Luna crash. Once again, if the run on Terra were induced by panic, the sentiment would decline, and disagreement would rise before the arbitrage mechanism broke down on May 7. To test this hypothesis, I collected — tweets about Terra from May 6, 2022, through May 22, 2022, and assigned each tweet a positive, negative, or neutral sentiment. I calculate the average and standard deviation of tweet sentiments per hour, the former representing the polarity of information on social media and the latter disagreement. Both measures were relatively constant before May 9, when sentiment declined and disagreement increased. Therefore, negative information did not spread until two days after the arbitrage mechanism had broken down.

The remainder of the paper is organized as follows. Section 2 reviews the related literature and discusses our contribution. Section 3 presents the institutional details. Section 4 presents the data. Section 5 outlines the mechanism to estimate arbitrage breakdown. Section 6 presents flight to safety. Section 7 outlines the empirical method to estimate adverse selection. Section 8 presents information transmission. Finally, Section 9 concludes.

2 Literature Review

This paper contributes to an emerging literature on stablecoins and brings unique evidence to studying run-like behavior in digital currencies by studying the crisis in stablecoin Terra. My primary contribution is to present evidence that identifies the causes of the crash. I concentrate on adverse selection and arbitrage limitations, constructing a timeline of events to demonstrate how the crash unfolded. With the rising significance of digital currencies as a forum for monetary transactions, a literature (Catalini and de Gortari, 2021) has emerged to study their design and stability .(Gorton and Zhang, 2023), (Makarov and Schoar, 2022) provide a comparative analysis of stablecoins and private money creation during the banking period of 19th century United States. (Liao and Caramichael, 2022) analyze the potential for stablecoins to broadly impact the banking system.

There is a growing literature that examines the Terra - LUNA crash. (Uhlig, 2022) propose a model for the Terra - LUNA crash and show how such a crash can unfold gradually. They argue that agents sell their UST coin, when the probability of an eventual suspension of convertibility to LUNA exceeds some convenience value of holding the UST coin. Suspension of convertibility happens, once the UST price has fallen sufficiently far. They find that the majority of the UST coin holders waited until the probability of suspension was rather high, before deciding to burn their holdings. Along the same lines, I study the evolution of arbitrage profits and define the structural stability of the algorithm by computing the parameters required to keep the algorithm stable. In my paper, arbitrage breakdown is analogous to suspension of convertibility. (Lyons and Viswanath-Natraj, 2023) argue that arbitrage by vault owners is a key stabilizing force, a finding consistent with my analysis of arbitrage stability in secondary markets in section 6. They analyze the stability mechanisms of Tether and Dai.

The cryptocurrency research most closely related to my paper focuses on Terra - LUNA crash. (Liu, Makarov, and Schoar, 2023) provide comprehensive details to the events that led to the run on Terra by using a detailed data from the Terra blockchain and trading data from exchanges. They find that the run on Terra was not the result of targeted market manipulation by a single entity, but rather stemmed from growing concerns about the sustainability of the system. Once a few large holders of UST adjusted their positions on May 7th, 2022, other large traders followed suit. My study references (Liu et al., 2023) to exploit the timeline of UST withdrawl from blockchain wallets to find the limits to arbitrage that triggered the arbitrage failure. My finding that a sudden shock to arbitrage stability on May 7, triggered by a massive UST withdrawal, led to depegging, aligns with their findings. Additionally, a significant increase in UST's adverse selection on May 7 supports concerns about the system's sustainability. They further document how different types of investors behaved during the run and exited from Terra. They show that large and more sophisticated investors were the first to run and used multiple avenues to exit UST and LUNA. In contrast, I show flight to safety dynamics from Terra to other stablecoins and net outflow from Terra by all investors.

(Ma, Zeng, and Zhang, 2023) analyze the run risk of USD-backed stablecoins and the design features that could affect the occurrence of run. They argue that stablecoins feature concentrated arbitrage. Their findings show that stablecoins with fewer arbitrageurs experience larger price deviations in secondary markets, which aligns with my findings. My results on the breakdown of arbitrage also support the seminal work by (e.g. (Shleifer and Vishny, 1997),(Gromb and Vayanos, 2002)) that imperfect arbitrage hurts price efficiency.

More generally, several other papers have explored the pegging mechanisms and stablecoin crash(Gorton, Klee, Ross, Ross, and Vardoulakis, 2022)(Liu et al., 2023)(d'Avernas, Maurin, and Vandeweyer, 2022)(Catalini and de Gortari, 2021)(Gorton and Zhang, 2023)(Liao and Caramichael, 2022). The pegging mechanism behind Terra and LUNA shares similarities with the extensive body of literature in international finance that explores the causes and consequences of currency crises in countries with fixed exchange rates.(Morris and Shin, 1998) frame a model with strategic interaction between the government and a group of speculators in the foreign exchange market. The speculators are uncertain about the behavior of other speculators, but their behavior depends non-trivially on what they believe they will do. Their model of self-fulfilling currency attacks is consistent with unique equilibrium model.

New public information about the fundamentals of securities can change their information sensitivity. (Foley-Fisher, Gorton, and Verani, 2020) study the adverse selection dynamics in CLO's by showing that adverse selection in AAA rated CLO tranches increased dramatically just after the COVID pandemic, January 1,2020. By capturing the adverse selection dynamics for UST and other cryptocurrencies, I contribute to the large empirical literature that deals with information asymmetry models of financial crisis: (Dang, Gorton, and Holmström, 2020) argue that a financial crisis can occur when the fundamentals are thought to have lost enough value to raise doubts among the traders that some may acquire private information. They also show that a financial crisis is a shift from information insensitive to information-sensitive short-term debt. (Brancati and Macchiavelli, 2019) provide empirical evidence for the shift in information production about the bank debt, from non-crisis times to crisis times. By focussing on banks' CDS spreads and the relations between median analysts' forecasts of banks' ROA and the dispersion of those forecasts, they find that more information is produced at the onset of the crisis. More analysts are assigned to cover banks, and the analysts produce significantly more precise information, measured by the standard deviation of banks' returns on assets (ROA). (Narayanan and Rhodes, 2022) examine information sensitivity and information production dynamics of AAA-rated Residential Mortgage Backed Securities during the financial crisis of 2007 - 2009. Along the same lines, I study information transmission and adverse selection dynamics and my results echo the findings of these papers.

(Gallagher, Schmidt, Timmermann, and Wermers, 2020) study information production, MMMFs redemptions and managers' rebalancing portfolios during Eurozone Crisis. They find that there was a significant selective information production about fund holdings and the flight to safety phenomenon was observed in the funds with the most sophisticated investors. My results on flight to safety dynamics are consistent with these results.

My paper also contributes to the growing literature on social media and run risks (Iyer and Puri, 2012) empirically show the role played by social media in worsening the crisis.(Cookson, Fox, Gil-Bazo, Imbet, and Schiller, 2023) investigate the role played by social media in the bank run on Silicon Valley Bank(SVB). Using comprehensive Twitter data, they provide comprehensive evidence that exposure to social media conversations about bank stocks amplify classical bank run risks.

I extend existing work on stablecoins in several ways. This study pushes beyond past work on the Terra - LUNA crash and pegging mechanisms. My work broadly contributes to the literature on financial crisis and bank runs. This study further contributes to the literature on currency crises and speculative attacks. A key difference is the design to defend the peg: the exchange rate of Terra coin with the US Dollar is defended with a second currency "in circulation", LUNA, rather than its own currency reserves.

3 Institutional Details

3.1 Stablecoins

Stablecoins are cryptocurrencies issued and redeemed for $1 on a blockchain, a decentralized digital ledger. They are designed to function as "money" in the cryptoverse by maintaining a stable peg of $1 (or some other reference currency). Stablecoins also trade in the secondary market on various crypto exchanges such as Binance, Coinbase, and Bitfinex, where arbitrageurs exploit the price difference between primary and secondary markets to restore the peg (see Fig 3). Their use as a medium of exchange has contributed to their meteoric rise. Table 1 lists the top 10 stablecoins based on their market capitalization as of 2023.

3.2 Stablecoin Backing

Stablecoins peg their prices to $1 by backing each token with at least $1 in US dollar-denominated assets as collateral. On-chain collateralized stablecoins are backed by assets that can be represented by tokens on a blockchain. Off-chain collateralized stablecoins are backed by bank deposits or other cash-like assets traded in the traditional financial system. Stablecoins are typically classified on the basis of their backing collateral or reserves.

3.2.1 Fiat-backed stablecoins

Fiat-backed stablecoins are stablecoins that are backed by reserves in fiat currency, such as the US dollar or the Euro. The reserves are typically held by a regulated institution, such as a bank, and are securely stored either in a bank vault or with a trusted financial custodian. These fiat funds are then tokenized on a blockchain 1:1 or at a slighly lower ratio (e.g., 1.2:1). The ability of these coins to maintain their peg is reliant upon the issuing entity maintaining significant reserves. Such stablecoins are typically issued by companies that operate centralized cryptocurrency service platforms, where the stablecoin functions as a medium of exchange. When holders wish to convert their stablecoins back to fiat, they initiate a transaction. Upon verification of the transaction, the corresponding stablecoins are removed from circulation to maintain the supply's parity with the fiat reserve. The issuer then releases the equivalent fiat funds to the holder. The two largest stablecoins by market capitalization, Tether (USDT) & USD Coin (USDC), fall in this category.

3.2.2 Crypto-backed stablecoins

Crypto-backed stablecoins are on-chain stablecoins backed by other cryptocurrencies where the holder can redeem the stablecoins for the collateralized cryptocurrencies on demand. Due to the high volatility of cryptocurrencies backing stablecoins, these stablecoins are typically over-collateralized and their stabilization mechanisms rely on continuous valuation of collateral and adjusting the reserve to reflect the collateral value. A prominent example of such a stablecoin is DAI which is maintained and regulated by MakerDAO on the Ethereum blockchain and backed by Ether (ETH).

3.2.3 Algorithmic stablecoins

Algorithmic stablecoins are a variant of crypto-backed stablecoins which rely on an algorithm to maintain their 1:1 price peg. When demand for the stablecoin increases and its price subsequently appreciates over $1.00, the algorithm issues (mints) new stablecoins to drop the price back to its pegged $1.00 price. Alternatively, when the price of the stablecoin falls below $1.00, the algorithm automatically destroys (burns) stablecoins using its reserve (often in another cryptocurrency) by sending them to a wallet from which they can't be retrieved. This reduction in supply is intended to decrease the number of coins in circulation, thereby driving the price back towards the peg. An example of an algorithmic stablecoin is Terra (UST).

Table 1: Market Cap and Backing Collateral of Top 10 Stablecoins (End of 2023)

Rank	Stablecoin	Market Cap (Billions USD, End 2023)	Backing Collateral
1	Tether (USDT)	$82.3	Fiat-backed
2	USD Coin (USDC)	$43.0	Fiat-backed
3	Binance USD (BUSD)	$11.0	Fiat-backed
4	Dai (DAI)	$7.5	Crypto-backed
5	TrueUSD (TUSD)	$3.1	Fiat-backed
6	Pax Dollar (USDP)	$1.0	Fiat-backed
7	Gemini Dollar (GUSD)	$0.6	Fiat-backed
8	HUSD	$0.52	Fiat-backed
9	sUSD (sUSD)	$0.42	Crypto-backed
10	Fei USD (FEI)	$0.33	Algorithmic

3.3 Market Structure

Stablecoin tokens are created ("minted") or redeemed ("burned") in the primary market with US dollar cash. To create a stablecoin token, an arbitrageur sends $1 to the issuer, who then sends a stablecoin token to the arbitrageur's crypto wallet. Conversely, to redeem a stablecoin token, the market participant sends a stablecoin token to the issuer's crypto wallet, and the issuer transfers $1, typically via bank transfer, to the participant's bank account. The primary market for stablecoins operates similarly to a money market fund in the traditional financial system (see Fig 3).

However, not all market participants can freely become arbitrageurs to engage in the redemption and creation of stablecoin tokens in the primary market. Access to primary markets varies among stablecoin issuers. For instance, USDC allows general businesses to register as arbitrageurs, while USDT requires a lengthy due-diligence process and imposes restrictions on the domicile of arbitrageurs. Additionally, USDT has a minimum transaction size of $100,000 and charges the greater of 0.1% or $1000 per redemption.

Most market participants trade existing stablecoins for fiat currencies in secondary markets. Crypto exchanges enable investors to deposit US dollars and trade them for stablecoins with other participants. The price of stablecoin tokens in the secondary market is driven by the demand from buyers and the supply from sellers. When there is a surge in stablecoin sales in the secondary market, prices will drop, but due to the closed-ended nature of stablecoins, these sales do not directly cause liquidations of reserve assets. Thus, the buying and selling of stablecoins in secondary markets are akin to trading ETF shares on competitive exchanges.

Selling pressure in the secondary market for stablecoins can impact the primary market through arbitrageurs. If investor selling pressure depresses stablecoin prices below \$1 in the secondary market, arbitrageurs can profit by purchasing stablecoin tokens at a discount and redeeming them one-for-one for \$1 with the issuer in the primary market, provided the issuer does not default. Similarly, if positive demand shocks cause stablecoins to trade above \$1 in secondary markets, arbitrageurs can profit by creating stablecoin tokens in the primary market and selling them at higher prices in the secondary market. Thus, the \$1 redemption value in primary markets pulls the trading price of stablecoins towards \$1 in secondary markets through arbitrage. This arbitrage process implies that investor selling pressure in secondary markets eventually triggers the liquidation of reserve assets by stablecoin issuers to meet arbitrageurs' \$1 redemption requests in cash.

3.4 Terra-Luna Design and Mechanism

TerraForm Labs (TFL) launched Terra in 2018 to smart contract blockchains like Ethereum. Its primary strategy was to create a wide range of blockchain-based applications and services to attract a stable user base and generate fees for cryptocurrency holders. At the heart of this ecosystem is LUNA, the collateral backing Terra, which derives its value from three key factors. First, LUNA holders receive a portion of Terra's transaction fees and block rewards by holding LUNA to support the blockchain network. Second, Luna serves as the means to access Terra's applications and drives transaction demand. Third, Luna's value could also be influenced by investor beliefs about the stability of the system.

Terra was backed by an algorithm that allowed an exchange of one unit of Terra for 1\$ worth of Luna and vice versa. The pegging mechanism relied on traders' exploiting arbitrage opportunities that arose whenever Terra deviated from its peg in either direction. Thus, when Terra is trading above \$1, users could buy LUNA, swap LUNA for Terra, which amounts to burning (destroying) LUNA and minting (creating) new Terra, and sell Terra at a premium above \$1, pocketing the difference as profit. In contrast, when Terra trades below $1, users could buy Terra, burn Terra to mint new LUNA, and then sell LUNA with a profit. In this system, arbitrage opportunities exist within secondary exchanges. Traditionally for other stablecoins, arbitrage profit is the difference between the discounted price paid in the secondary exchange and the $1 redemption value obtained on the blockchain.

3.5 The History of the Terra-Luna System

The Terra network was developed by Terraform Labs (TFL) and was founded by Do Kwon and Daniel Shin in 2018, both entrepreneurs from South Korea. The founders established Terraform Labs with the vision of creating a decentralized stablecoin ecosystem. The project aimed to create an efficient, scalable, and decentralized payment system that could be used worldwide, offering low transaction fees and fast processing times. In 2018, TFL successfully secured $32 million in seed funding from multiple venture funds and prominent cryptocurrency exchanges, such as Binance and Huobi. The network officially launched in April 2019.

Terra initially gained traction through partnerships with various e-commerce platforms, particularly in South Korea. The network's native stablecoin, TerraUSD (UST), and its governance token, LUNA, played crucial roles in its growth. The total UST and LUNA volumes during the pre-crash period are $360 and $217 billion, respectively(Liu et al., 2023).Initially, Terra's market capitalization was relatively small, beginning in the low millions as the ecosystem and user base were developed. By May 2022, Terra's market cap had soared to over $18 billion [1]. LUNA's price surged as UST's adoption increased, making it one of the top ten cryptocurrencies by market capitalization.

One of the main attractions on Terra was the borrowing and lending protocol, Anchor, which provided heavily subsidized deposit rates of 20% until the beginning

[1]https://coinmarketcap.com/currencies/terra-luna/

of the crash on May 7, 2022[2], and drew many users to the platform. Anchor was the most significant protocol within the UST network, comprising 46% of the total network volume. (Liu et al., 2023) show that the subsidies provided by Anchor remained the primary attraction for Terra users. In response to widespread concerns about the unsustainability of the 20% Anchor rate, several proposals to reduce it emerged in March 2022. Anchor Proposition 20, approved on March 23 and enacted on May 1, 2022, linked the deposit rate to the monthly change in the yield reserve. Under this proposition, the rate would increase if the yield reserve grew and decrease if it declined, with a maximum monthly adjustment of 1.5%.

On May 7, the first signs of a run on TerraUSD (UST) were reported in various industry reports and on social media. Large withdrawals from Anchor Protocol triggered a sell-off in UST. Panic ensued as UST began to lose its peg to the US dollar, dropping below $0.98. Two addresses, referred to as wallet A and wallet B, initiated significant withdrawals from Anchor. Wallet A withdrew 45 million UST around 5:00 AM UTC and transferred the funds to Binance. Following this, wallet B withdrew 175 million UST around noon, sending the funds to Ethereum via the Wormhole bridge. Later, wallet A withdrew another 35 million UST around 5:00 PM and an additional 20 million UST at 8:30 PM and 9:30 PM, all of which were sent to Binance. At 9:44 PM, Terraform Labs (TFL) removed 150 million UST from the UST-3Crv pool to transfer it to a new UST-4Crv pool. Finally, at 9:48 PM, wallet A withdrew the last 85 million UST and sent them to Curve. The algorithmic mechanism failed to stabilize UST, causing a sharp decline in UST's value. On May 9, 2022, the situation worsened as more UST holders rushed to exit their positions, leading to further de-pegging and the eventual collapse of UST.

After the crash, the Terra community and Terraform Labs decided to fork the Terra blockchain to create a new version. The original chain retained the name

[2]https://docs.anchorprotocol.com/anchor-2/protocol/anchor-governance/modify-market-parameters

Terra (LUNA), while the new chain was rebranded as Terra Classic (LUNC). The new Terra chain moved forward without the algorithmic stablecoin, TerraUSD (UST), to avoid the issues that led to the previous collapse and focus on developing a more robust ecosystem. Regulators in several countries launched investigations into Terraform Labs and Do Kwon's role in the collapse. South Korean authorities issued an arrest warrant for him, citing violations of capital markets law and fraud. In September 2022, Interpol issued a Red Notice for Do Kwon, effectively making him a fugitive wanted by law enforcement agencies globally. This notice was issued at the request of South Korean prosecutors. On March 23, 2023, Do Kwon was arrested in Montenegro while attempting to board a flight using allegedly falsified documents. He remains in custody there as legal proceedings continue. Both South Korea and the U.S. have requested his extradition, and he is likely to face trial in one or both countries, but the outcome of these extradition requests has yet to be determined.

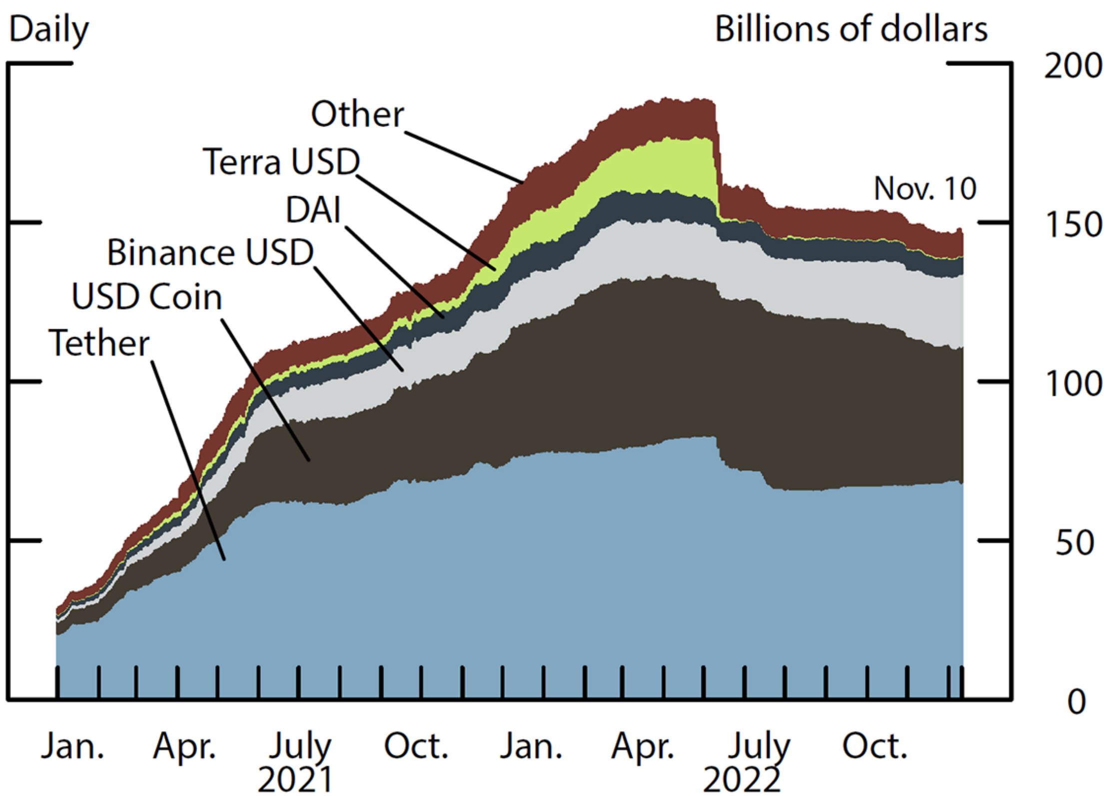

Figure 1: Stablecoins' Market Capitalization, Source: FRB

Table 2: Top 10 Crypto Exchanges by Trading Volume in 2024

Rank	Exchange	Trading Volume (Jan 2024)	Trading Volume (Feb 2024)	Trading Volume (Mar 2024)
1	Binance	$2.1 trillion	$2.3 trillion	$2.5 trillion
2	Coinbase	$1.8 trillion	$1.9 trillion	$2.0 trillion
3	Kraken	$1.2 trillion	$1.3 trillion	$1.4 trillion
4	Bitfinex	$1.1 trillion	$1.2 trillion	$1.3 trillion
5	OKEx	$0.9 trillion	$1.0 trillion	$1.1 trillion
6	Huobi Global	$0.8 trillion	$0.9 trillion	$1.0 trillion
7	Bitstamp	$0.7 trillion	$0.75 trillion	$0.8 trillion
8	Bittrex	$0.6 trillion	$0.65 trillion	$0.7 trillion
9	KuCoin	$0.55 trillion	$0.6 trillion	$0.65 trillion
10	Gate.io	$0.5 trillion	$0.55 trillion	$0.6 trillion

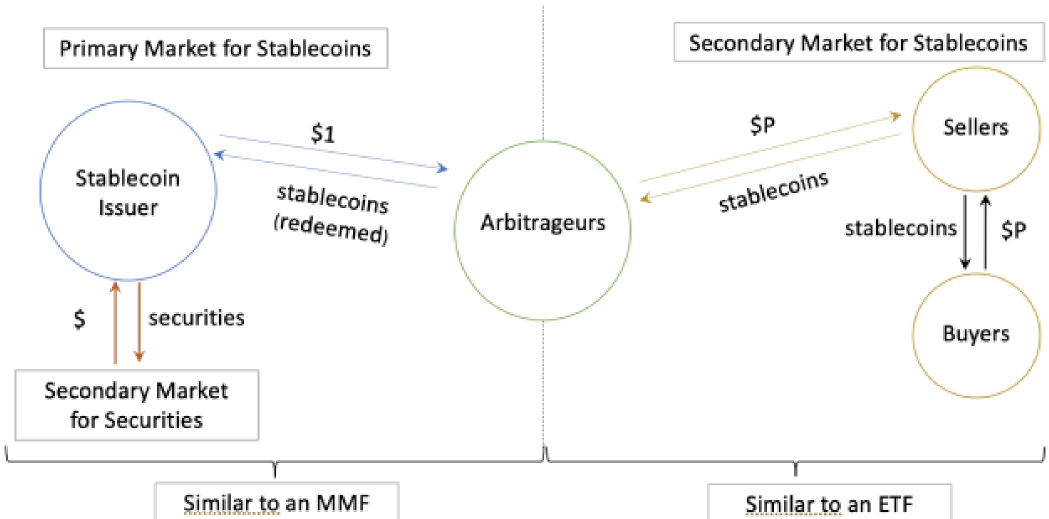

Figure 2: Arbitrage Mechanism for Fiat Backed Stablecoins
This figure, sourced from (Ma et al., 2023), illustrates the critical role of arbitrageurs in maintaining the stablecoin's price at 1 USD. The primary market represents the blockchain, where stablecoins can be redeemed for 1 USD. In contrast, secondary markets refer to cryptocurrency exchanges such as Binance, Coinbase, and Bitfinex, among others, where price deviations from 1 USD are often observed. Arbitrageurs exploit the price difference, $1 - P$, between the primary and secondary exchanges to maintain the stablecoin's price at its peg.

4 Data

The dataset includes trading data from Binance for the cryptocurrency pairs analyzed in this paper, covering April 1, 2022, to May 12, 2022. If unavailable on Binance, data is sourced from other major crypto exchanges. Tweet data is collected from Twitter, Reddit, and Instagram between May 6, 2022, and May 13, 2022.

4.1 Trading Data

I use tick-by-tick trading data to compute a measure for arbitrage functioning every hour for the cryptocurrencies studied. The trade data includes trade direction, allowing me to compute the signed volume of each trade.I define arbitrage profits at time t as $|1 - P_{\text{UST},t}|$ where $P_{\text{UST},t}$ is the timestamped price of UST. Using these arbitrage profits, I estimate β, a proxy of algorithm functioning, using the autoregressive (AR1) model. I compute β every hour by running AR1 regressions on tick-by-tick trade. Similarly, I compute daily values of β from May 4 to May 12. I interpret an increase in β as an algorithm malfunction. I use May 7 as the start of the crisis. This timing is consistent with a significant jump in the value of β (the first signal of arbitrage failure) and adverse selection. Prior to the crisis, β values are remarkably low, and there is virtually no adverse selection (no compensation from market makers for potential adverse selection risks): all stablecoins are perceived as equally safe.

I compute Net Dollar Outflow using the variables in the data. The **UST to Crypto$_i$ Sell Volume** (UST_Crypto$_i$_Sell_Vol) represents the total volume of UST sold to acquire Crypto$_i$. The **UST to Crypto$_i$ Buy Volume** (UST_Crypto$_i$_Buy_Vol) represents the total volume of UST bought by selling Crypto$_i$.

Using the same data, I compute the adverse selection measure every hour for the cryptocurrencies studied. To normalize the adverse selection for all cryptocurrencies, I divide the hourly adverse selection measure by the average hourly price of each cryptocurrency. I use this normalized measure of adverse selection to analyze stablecoins before, during, and after the crash, and to compare the adverse selection of UST with other cryptocurrencies.

4.2 Tweet Data

I collect 5000 tweets per hour from 6 May 2022 until 13 May 2022 using # LUNA. I compute tweet polarity every hour using Natural Language Processing. I compute the standard deviation of tweet polarity to capture the extent of investor disagreement every hour.

5 Arbitrage Breakdown

The Terra-LUNA algorithm was designed to maintain Terra's peg to the US dollar resulting in arbitrage opportunities. As Terra's price deviation from its peg widened, the profitability of engaging in these arbitrage activities increased. Investors found it increasingly lucrative to swap Terra for LUNA or vice versa, depending on the price of Terra.

The arbitrage process functions effectively when large arbitrage profits (indicated by $|1 - P_{\text{UST},t-1}|$) incentivize arbitrageurs to push the price of Terra (UST) back to one dollar. Consequently, in the subsequent trade, the arbitrage profits should decrease, bringing $|1 - P_{\text{UST},t}|$ closer to zero. To quantify this relationship, I estimate the coefficient β hourly using the AR1 model as described:

$$|1 - P_{\text{UST},t}| = \alpha + \beta \times |1 - P_{\text{UST},t-1}| + \epsilon_t \tag{1}$$

where $P_{\text{UST},t}$ represents the price of UST at time t, and ϵ_t denotes the error term. The coefficient β in this model measures the persistence of price deviations from the peg over time. I use tick-by-tick data from Binance and run the regression equation 1 using all the timestamped trades as data points every hour. This AR1 model captures the dynamic adjustment of UST prices in response to arbitrage activities, validating the effectiveness of the arbitrage mechanism in stabilizing Terra's price.

Arbitrage fails when the UST price remains unchanged from the previous trade at $t-1$ or when the coefficient β approaches 1:

$$\text{Arbitrage fails if } P_{\text{UST},t} = P_{\text{UST},t-1} \text{ or } \beta \to 1.$$

The long-term value of arbitrage profit $|1 - P_{\text{UST},t}| = |1 - P_{\text{UST},t-1}|$ is given by $\frac{\alpha}{1-\beta}$ as in the AR(1) process. As β approaches 1, the system becomes unstable, causing persistent deviations from the peg. A high value of β means that a large portion of the price deviation from the peg persists over time, indicating that the arbitrage mechanism is less effective in correcting the price. Low β indicates low information asymmetry and efficient market correction.

Figure 3 illustrates that β starts at a moderate level of around 0.2 on May 4 and remains relatively stable until May 6. β starts to consistently increase from May 7, 12 PM and reaches a value of 1 at 5 PM (Figure 4). This implies that on May 7, post noon, arbitrageurs were unable to restore UST to its peg, marking the first observed failure of the mechanism at 5 PM.

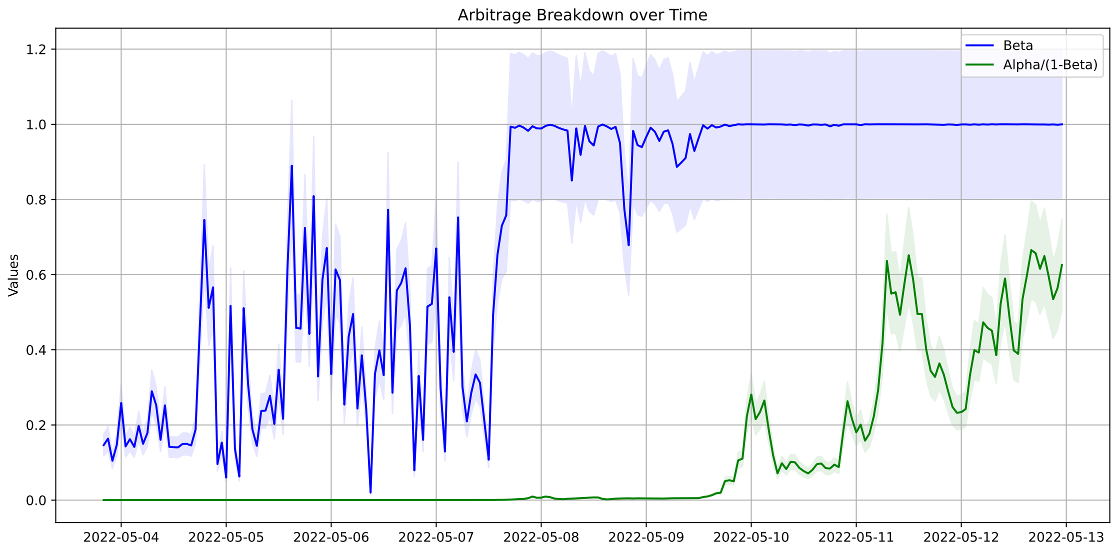

Figure 3: Hourly β and $\frac{\alpha}{1-\beta}$ for Terra (UST) This figure shows the hourly values of coefficient β and the stable value of arbitrage profits in our sample from May 4, 2022 to May 12, 2022. The solid blue and green lines display, respectively, the values of coefficient β and arbitrage 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 on the system. The plot shows the values of and LTV ($\frac{\alpha}{1-\beta}$) every hour, with key withdrawal events marked on the timeline. At the start of May 7, both and LTV exhibit stable behavior, indicating that the arbitrage mechanism is functioning effectively. A large withdrawal of 175 million USD 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 LTV, indicating increased market risk and increased arbitrage profitability. Importantly, this breakdown occurs before any noticeable increase in the adverse selection measure (Figure 8). This sequence of events demonstrates that the arbitrage breakdown is not caused by adverse selection. Instead, it suggests that the breakdown could have triggered the subsequent increase in adverse selection.

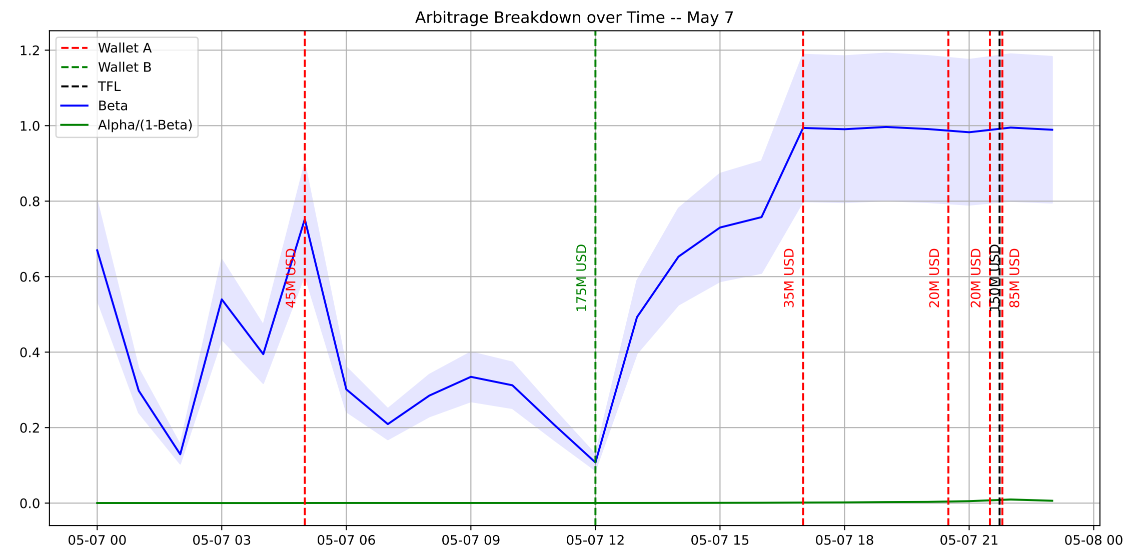

Figure 4: May 7 β and $\frac{\alpha}{1-\beta}$ for Terra (UST) This figure shows the values of coefficient β and the stable value of arbitrage profits for May 7, 2022. The solid blue and green lines display the values of coefficient β and arbitrage profits, $\frac{\alpha}{1-\beta}$. This figure shows major UST withdrawals and the evolution of β around the withdrawals. 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, and the number of observations (daily trades) as presented in Table 3. 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 3: Estimates for Arbitrage Breakdown

In this table, I estimate the value of β and the stable value of arbitrage profits (LTV Arbitrage Profits), $\frac{\alpha}{1-\beta}$, using the Autoregressive model of order 1. The dependent variable is the value of arbitrage profits, computed daily as $|1 - P_{\text{UST},t}|$. The independent variable is the daily value of arbitrage profits at time $t-1$. The parameter, $|\beta|$ must be less than 1 for the model to be stationary. The arbitrage mechanism fails when the estimated value of β tends to 1. The table shows the Beta values and Long Term Value of Arbitrage Profits from May 4, 2022, to May 12, 2022. The values of β are rounded to 3 decimals and the values of $LTV\,ArbitrageProfits$ are rounded to 2 decimals.

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
$LTV\ Arbitrage\ Profits(cents)$	0.01	0.01	0.02	0.51	0.52	10.79	16.46	45.95	50.91
Observations	16123	19768	15764	41985	77072	190426	751686	1744320	1566119

6 Flight to Safety

In this section, I show that the flight-to-safety dynamics occur during run episodes in the stablecoin market similar to the dynamics observed in the run on Money Market Mutual Funds (MMMF) (2008 and 2020). During a stress event, investors identify a driver of risk and run away from stablecoins that are more exposed to such risk toward relatively safer stablecoins or similar vehicles.

To quantify these movements, I compute net dollar outflow from UST to Crypto_i using equation 2

$$\text{Net Dollar Outflow from UST} = \left(\sum_i \left(\text{UST\_Crypto}_i\text{\_Sell\_Vol} \times \frac{\text{UST Avg Price}}{\text{Crypto}_i \text{ Avg Price}}\right)\right) - \left(\sum_i \left(\text{UST\_Crypto}_i\text{\_Buy\_Vol} \times \frac{\text{Crypto}_i \text{ Avg Price}}{\text{UST Avg Price}}\right)\right) \times \frac{1}{10^6} \tag{2}$$

Equation (2) captures the financial dynamics between UST and various cryptocurrencies by calculating the net flow of dollars resulting from trades between these assets. The **UST to Crypto_i Sell Volume** ($\text{UST\_Crypto}_i\text{\_Sell\_Vol}$) represents the total volume of UST sold to acquire Crypto_i, indicating the selling pressure on UST as traders exchange it for Crypto_i. Conversely, the **UST to Crypto_i Buy Volume** ($\text{UST\_Crypto}_i\text{\_Buy\_Vol}$) reflects the total volume of Crypto_i purchased using UST, highlighting the buying pressure on Crypto_i from UST holders. The **Price Adjustment Ratios**, $\left(\frac{\text{UST Avg Price}}{\text{Crypto}_i \text{ Avg Price}}\right)$ and $\left(\frac{\text{Crypto}_i \text{ Avg Price}}{\text{UST Avg Price}}\right)$, adjust the traded volumes to account for price differences between UST and Crypto_i. These ratios ensure that the trade volumes are comparable in dollar terms, accurately reflecting the financial impact of these transactions. Finally, the **Normalization Factor** $\left(\frac{1}{10^6}\right)$ converts the calculated outflow into millions of dollars.

Figure 5 illustrates that massive flows out of Terra happened on May 9 around 6 PM. The net outflow from UST is computed by aggregating the outflows to

LUNA, BTC, USD, USDC, and USDT. During the early days (May 1 to May 6), the net outflows from UST remained relatively low. This period indicates a state of equilibrium where the inflow and outflow of UST were balanced, suggesting that the market was stable. A noticeable outflow starting on May 7 is observed, however, the outflows are not pronounced.

Between May 8 and May 10, the net outflows from UST surged dramatically, reaching hundreds of millions of dollars. This sharp increase in outflows reflects a loss of confidence among investors following the breakdown of the arbitrage mechanism. As the peg stability of UST weakened, investors started moving their funds out of UST.

On May 11, a critical shift occurred in the market dynamics. Initially, there was a brief period of net inflow into UST as some investors attempted to take advantage of increased arbitrage opportunities or believed that the peg could be restored. This momentary optimism, however, was short-lived. Following this brief inflow, the market experienced massive net outflows. The net outflow plot shows a dramatic increase, with outflows exceeding 300 million USD, far surpassing any previous levels. This marked a definitive market run as confidence in Terra's system eroded.

The timeline of net outflows suggests that massive outflows from Terra's UST occurred after May 11, however, arbitrage broke down completely on May 8. Therefore, the breakdown in unlikely to be a result of an informational event.

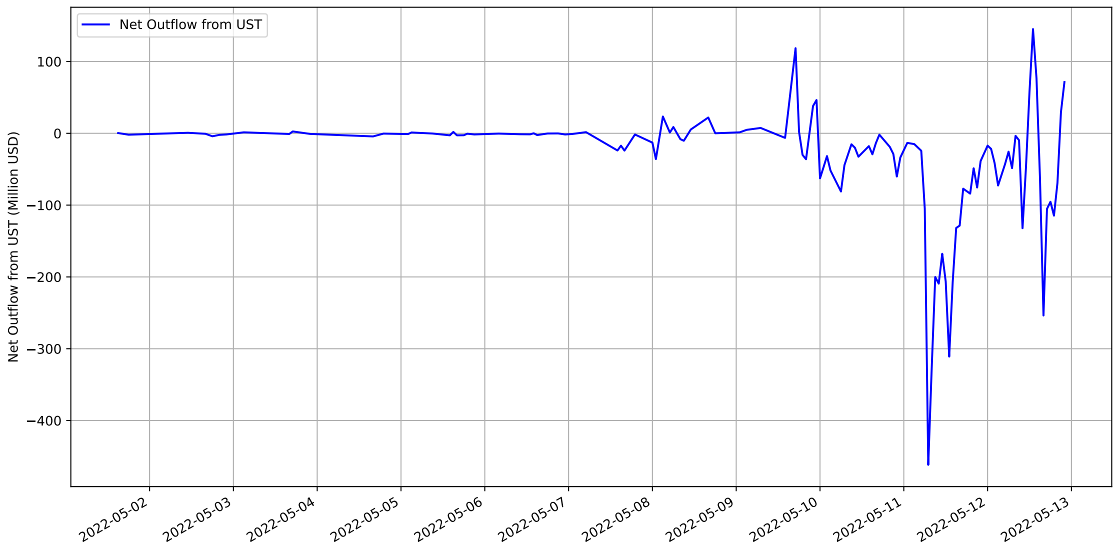

Figure 5: Net outflows from Terra around collapse : This figure shows the time series of net outflows from Terra during the period: May 1, 2022, to May 13, 2022. Significant outflows were observed around the crash. However, following the Terra crash post 5 PM on May 9, there was a significant net inflow to Terra.

The data, spanning from May 1 to May 13, is further analyzed to capture the critical period of Terra's market shifts. Figure 6 illustrates the distribution of outflows from Terra in terms of different coins both before and after noon on May 7. This provides insight into the shifts in market behavior as the arbitrage breakdown unfolded. The stark contrast between pre- and post-noon May 7 outflows demonstrates a clear shift in investor behavior. Before the arbitrage breakdown at noon on May 7, outflows were more diversified among USDT, USD, and USDC, reflecting a balanced risk perception. USDT was still the dominant asset in outflows, accounting for 70.9%, amounting to 11.43 million USD, of the total outflows. Investors did not yet face immediate panic, allowing for a more measured approach in choosing assets for withdrawal.

After noon on May 7, USDT became the dominant asset in outflows, accounting for 96.2%, amounting to 4927 million USD, of the total outflows. This shift indicates a strong preference for USDT. USDT's widespread acceptance on major exchanges and its high liquidity made it an attractive option for investors looking to quickly exit Terra's ecosystem.

Before noon on May 7, 22.5% of the outflows, equivalent to 3.64 million USD, went into fiat USD, indicating that a significant number of investors were opting to convert their crypto assets into traditional currency amid early signs of market instability. After noon on May 7, 3.7% of the outflows, amounting to 189.64 million USD, were directed into fiat USD. Although this is a significant amount in absolute terms, the percentage of outflows into fiat USD decreased considerably as investors shifted their preference toward stablecoins like USDT within the crypto ecosystem. The crash did not have major effects on the traditional financial sector because only a small portion of the capital in the Terra ecosystem flowed into USD. This limited conversion to traditional currency meant that the majority of the liquidity crisis was contained within the crypto space, minimizing potential spillover effects on conventional financial markets.

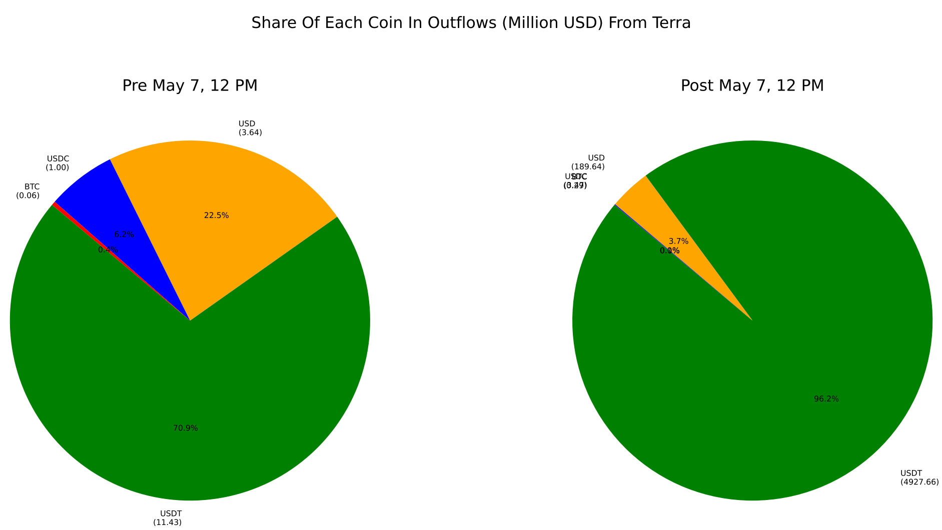

Figure 6: Net outflows from Terra around collapse : This figure shows the net outflows from Terra during the period: May 1, 2022, to May 13, 2022. USDT shows the highest dollar inflows from Terra.

7 Estimation of Adverse Selection

I follow Glosten and Harris (1988) to estimate the adverse selection component of bid-ask spreads. Accordingly, I denote the unobserved true price of the asset by m_t. It represents the price of the asset assuming a fully competitive market maker and no inventory costs or clearing fees. Innovations in m_t, therefore, result from public news or information arrival through order flows. More formally,

$$m_t - m_{t-1} = e_t + Q_t Z_t, \tag{3}$$

with $Q_t = 1(-1)$ for a buyer- (seller-) initiated trade. In Equation 3, e_t represents the impact of public news, and Z_t is the market maker's compensation for bearing the adverse selection risk. Hence, Z_t is the adverse selection component of the bid-ask spread. The observed prices reflect factors such as the market maker's monopoly power, inventory costs, and clearing fees. To account for these factors, Glosten and Harris (1988) include a second component in their specification, called the *transitory component*. Formally,

$$P_t = m_t + Q_t C_t, \tag{4}$$

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

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

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

Substituting these into Equation 4 and taking the first difference gives

$$P_t - P_{t-1} = c_0(Q_t - Q_{t-1}) + c_1(Q_t V_t - Q_{t-1} V_{t-1}) + z_0 Q_t + z_1 Q_t V_t + e_t. \tag{5}$$

Using trade data from Binance, I estimate the parameters in Equation 5 for every trading hour from May 5, 2022 to May 13 , 2022. I observe the trade direction, price, and volume. Therefore, I observe all the variables in Equation 5 and can estimate the parameters, c_0, c_1, z_0 and z_1 through OLS regressions.[3] Using the estimated parameters, $\hat{c}_0, \hat{c}_1, \hat{z}_0$, and $\hat{z}_1$, I calculate $\hat{C}_t$ and $\hat{Z}_t$ for each trade in the data. My measure of the adverse selection component of the bid-ask spread for the trading hour h is

$$AdverseSelection_h = \frac{\sum_{t \in h} \hat{Z}_t}{\mu_{t \in h}(Price_t)}$$

7.1 Evolution of Adverse Selection

Table 4 shows the magnitude of $AdverseSelection_h$ across different cryptocurrencies from May 1, 2022, to May 13, 2022. The measure of adverse selection represents the market maker's compensation for risk per dollar of traded asset. For Terra, an adverse selection measure of 0.0034 before May 7 indicates that market makers required only 0.000034 cents per dollar traded to cover adverse selection risk, highlighting Terra's relative stability at that time.

Before the crisis, the adverse selection measures for the stablecoins—Terra, USDC, and USDT—were very low, at 0.0034 bps, 0.0019 bps, and 0.0182 bps, respectively. This reflects the minimal risk of adverse selection typically associated with stablecoins. In comparison, non-stablecoins like LUNA and Bitcoin had higher adverse selection measures of 1.2779 bps and 0.2015 bps, respectively. These higher

[3] Contrary to Glosten and Harris (1988), I can observe the trade direction (Q_t) in my data. Furthermore, the rounding error in prices is negligible in my setup, given the minuscule tick size on Binance (0.00000010 for LUNA/UST pair). Therefore, I do not have to resort to the maximum likelihood method to estimate the parameters. Glosten and Harris (1988) mention this possibility, too.

measures reflect greater perceived risk, which demands higher compensation from market makers for potential adverse selection risks.

During the crisis, I observe a sharp increase in the adverse selection measure for Terra, which rose to 0.1019 bps, indicating significant instability and heightened risk perception. This reflects the severe impact of the market disruption on Terra, leading to increased uncertainty and adverse selection issues. Conversely, USDC and USDT maintained relatively low adverse selection measures of 0.0015 bps and 0.0250 bps, respectively, even during the crisis. This suggests that these stablecoins were less affected by the crisis, likely due to their robust pegging mechanisms and market confidence. For non-stablecoins, LUNA experienced a slight increase to 1.3906 bps, and Bitcoin increased to 0.2557 bps, indicating a general rise in risk perception during the crisis, although not as pronounced as Terra's spike.

Post-crisis, the adverse selection measure for Terra further increased to 1.1724 bps, illustrating ongoing instability and market concerns surrounding Terra. This sustained high measure suggests that Terra faced prolonged adverse selection issues, highlighting the lasting impact of the crisis on its market perception. Meanwhile, USDC and USDT continued to demonstrate resilience, with measures of 0.0062 bps and 0.0652 bps, respectively, in the post-crisis period. These figures indicate that these stablecoins maintained their stability, reassuring market participants of their reliability and no spillovers from the crash. In contrast, LUNA's adverse selection measure surged dramatically to 13.2527 bps, reflecting extreme volatility and risk. Bitcoin also saw an increase to 0.6860 bps, indicating heightened risk perception in the broader cryptocurrency market.

Table 4: Summary Statistics for Adverse Selection Measures

In this table, I estimate the value of adverse selection for various stablecoins using (Glosten and Harris, 1988). The table shows the evolution of adverse selection for various cryptocurrencies. This study primarily compares the adverse selection in Terra with other coins over different periods. Adverse selection is computed in bps.

Cryptocurrency	**Pre May 7**	**May 7 - May 9**	**Post May 9**
Terra	0.0034	0.1019	1.1724
USDC	0.0019	0.0015	0.0062
USDT	0.0182	0.0250	0.0652
LUNA	1.2779	1.3906	13.2527
Bitcoin	0.2015	0.2557	0.6860

Figure 7 illustrates how the adverse selection measure evolves, particularly focusing on significant fluctuations corresponding to market events. The first major increase in the adverse selection measure occurs on May 7. The measure rises sharply from a baseline of around 0.00000 to 0.00007, indicating that market makers began demanding higher compensation for taking on additional adverse selection risk.

Another significant spike occurs on May 9, extending into May 10. This second wave sees the adverse selection measure reaching its peak, surpassing 0.00012. This substantial increase highlights a renewed surge in market risk perception, possibly due to continued large withdrawals or further destabilizing events affecting Terra's ecosystem. The peak suggests heightened caution among market makers, who now require even greater compensation for the perceived risk. Post-crisis, the adverse selection measure does not revert to pre-crisis levels. Instead, it stabilizes at a higher baseline, suggesting lasting impacts on market perception and confidence.

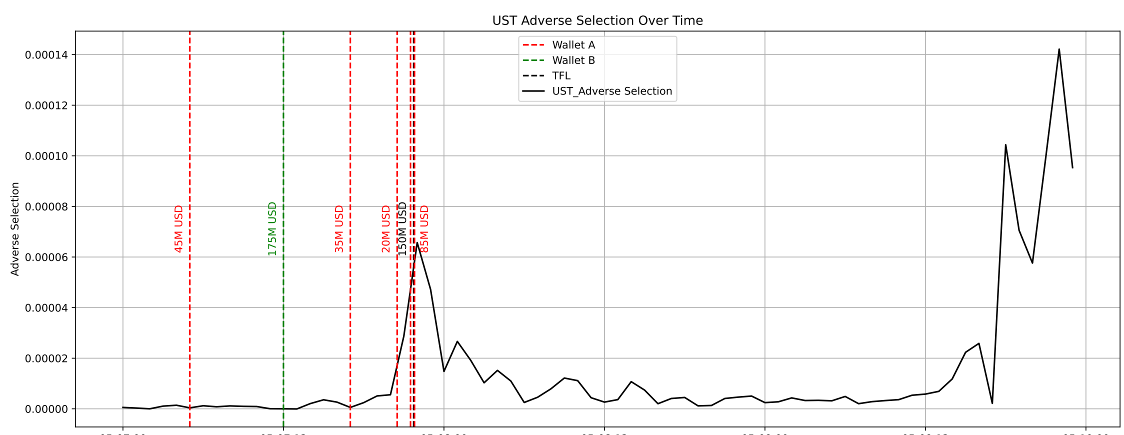

Figure 7: Intraday Adverse Selection – Terra. The figure shows the timeline of major UST withdrawals and the evolution of adverse selection in our sample from May 7, 2022 to May 9, 2022. A significant increase in adverse selection on May 7 and May 9 suggests that the withdrawals might have played a role in the increase. The data is from Binance.

Figure 8 illustrates how large withdrawals from specific wallets impacted the adverse selection, emphasizing the relationship between market actions and perceived risk by market makers. In the early hours of May 7, the adverse selection measure remains low, indicating a stable market environment. The measure fluctuates around a relatively low baseline, suggesting confidence in the stability of UST.

A significant event occurs when a large withdrawal of 175 million USD takes place. This withdrawal is marked on the graph and coincides with a notable change in the adverse selection measure. Following the 175 million USD withdrawal, there is an immediate and pronounced increase in the adverse selection measure. This spike indicates that market makers suddenly perceive higher risks associated with trading UST, prompting them to demand greater compensation for adverse selection. The measure rises sharply from its low baseline, demonstrating how significant financial outflows can destabilize market confidence. After the initial spike, the adverse selection measure does not return to its earlier stable levels, and remains elevated, reflecting ongoing uncertainty and risk perception.

Overall, the figure shows that the initial stability and subsequent spike in adverse selection following the 175 million USD withdrawal illustrate how quickly market dynamics can shift due to large financial outflows.

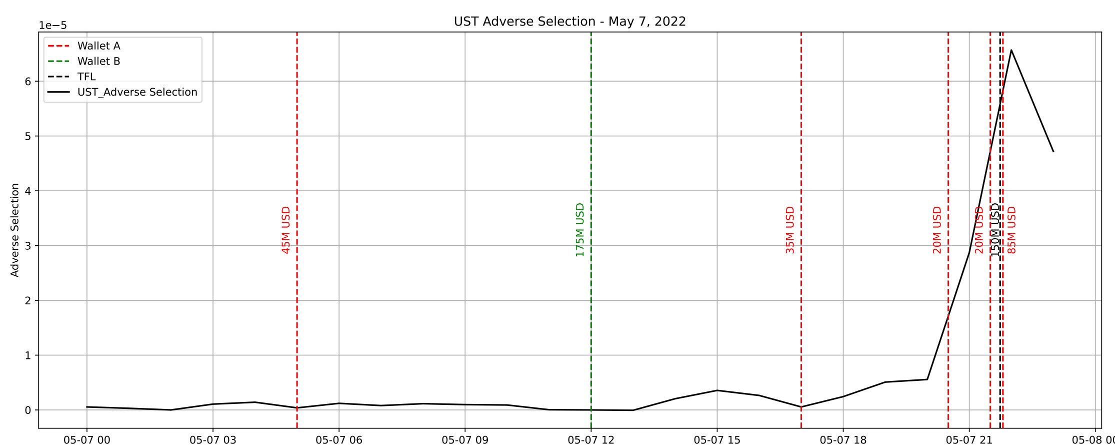

Figure 8: Adverse Selection May 7 – Terra. The figure shows the timeline of major UST withdrawals and the evolution of adverse selection for May 7, 2022. A significant increase in adverse selection on May 7 suggests that the withdrawals might have played a role in the increase. The data is from Binance.

Figure 9 provides a visual representation of the adverse selection for both LUNA and UST over the period from May 4, 2022, to May 13, 2022. The plot clearly shows how both LUNA and UST experienced significant increases in adverse selection relative to their price, particularly around May 7 to May 10. The simultaneous spikes and fluctuations in the adverse selection for both LUNA and UST suggest that these two assets may have reacted similarly to external market pressures.

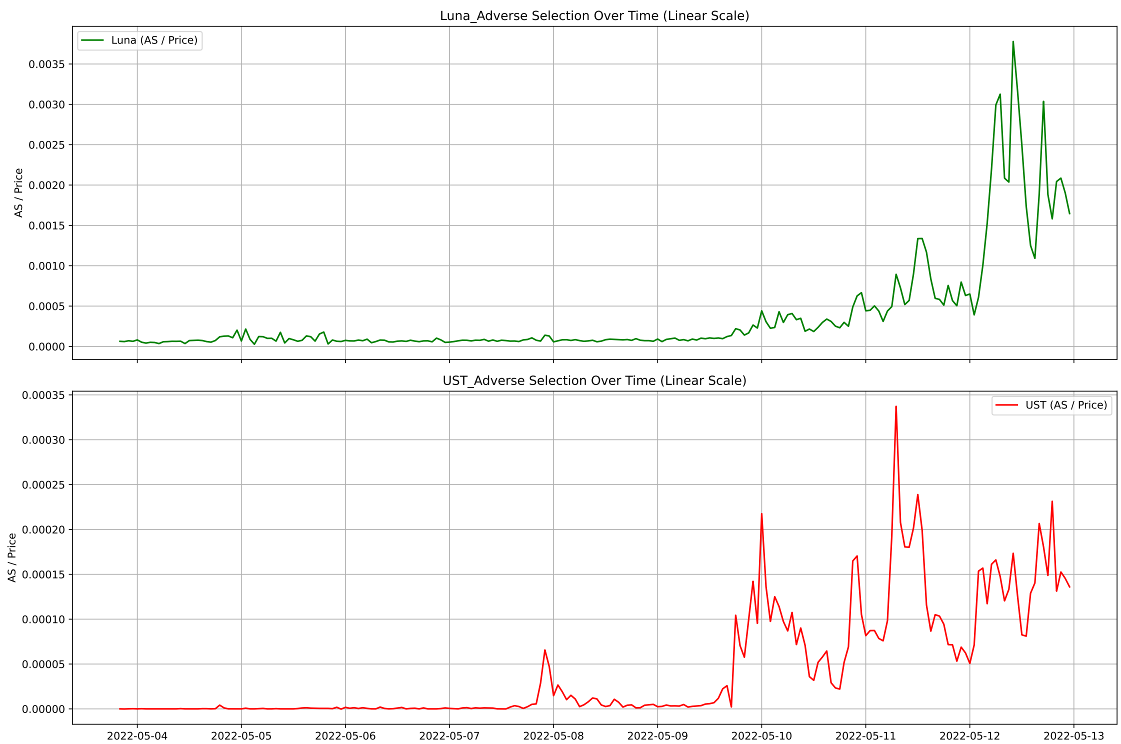

Figure 9: Adverse Selection Component of Bid-Ask Spread for Terra and LUNA. This figure shows the evolution of adverse selection for Terra and LUNA in our May 5 – May 13 sample. The solid green and red lines display, respectively, the adverse selection for LUNA and Terra in the sample. The data for computing adverse selection is from Binance (UST/USD) and Bitfinex (LUNA/USD).

7.2 Evolution of adverse selection in other cryptocurrencies

Bitcoin, the largest cryptocurrency by market cap, is often seen as representative of the crypto market. USDC has historically maintained a value close to $1 USD, making its adverse selection a benchmark for all stablecoins. Figure 10 provides a visual representation of the adverse selection for both LUNA and UST over the period from May 5, 2022, to May 13, 2022. The plot illustrates how the market's perception of risk changes over time for these two different types of cryptocurrencies: a highly volatile cryptocurrency (BTC) and a stablecoin (USDC). While Bitcoin shows notable fluctuations in its adverse selection measure, USDC maintains a consistent value close to 0.

The variations in Bitcoin's adverse selection further suggest that it faced some spillover effects from the UST crash, as market participants adjusted their risk assessments in response to the broader market instability. In contrast, USDC's stable nature buffered it from such fluctuations, highlighting its role as a safe haven during periods of market uncertainty.

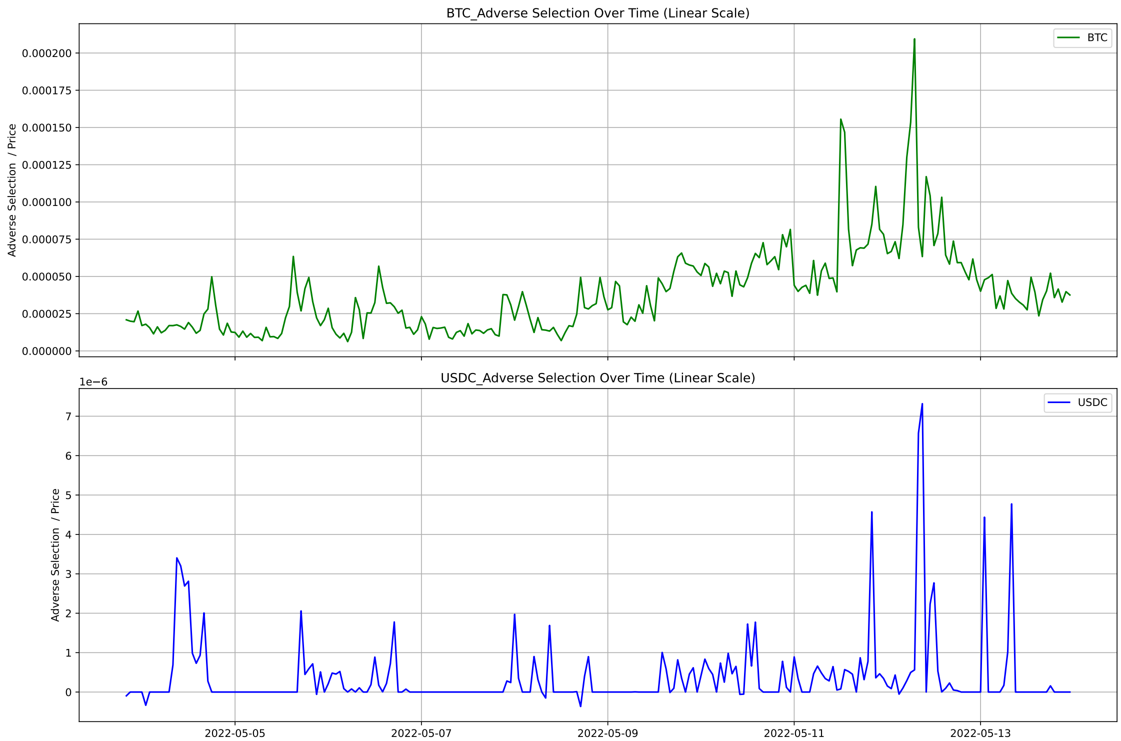

Figure 10: Adverse Selection Component of Bid-Ask Spread for Bitcoin and USDC This figure shows the evolution of adverse selection for Bitcoin (BTC) and stablecoin USDC in our May 5 – May 13 sample. The solid green and blue lines display, respectively, the adverse selection for BTC and USDC in the sample. The data is from Binance.

8 Information Transmission

(Iyer and Puri, 2012) argue that social networks influence investors' propensity to run. Depositors observe the other depositors' actions only if connected by the network. Using data from social media such as Twitter, Reddit and Instagram, I collect 5000 tweets per hour for LUNA from 6 May 2022 until 13 May 2022 to study the impact of information transmission on investor-run behavior. Remember that UST and LUNA are interlinked and LUNA serves as collateral for UST. Investors' tweets at time $t-1$ influence the market dynamics at time t.

I use TextBlob, a Python library specifically designed for social media for natural language processing (NLP), to score the sentiment of each tweet. TextBlob can perform various NLP tasks such as part-of-speech tagging, sentiment analysis, noun phrase extraction, translation, and classification. I compute the negative or positive tweet volume by analyzing tweet sentiment and calculate the polarity of tweets by providing a raw sentiment score ranging from -1 to 1. Negative, zero, and positive polarity scores represent negative, neutral, and positive sentiments respectively.

Figure 11 captures the sentiment dynamics surrounding LUNA from May 6 to May 12, 2022. Initially, from May 6 to the morning of May 9, sentiment remained relatively stable and slightly positive, with polarity values between 0.10 and 0.12. However, the sentiment took a dramatic downturn following the public realization of the crash on the afternoon of May 9. The polarity fell sharply, reaching lows of approximately 0.02 by May 10 and remaining low through May 12. This drastic decline indicates that investors rapidly lost confidence in LUNA and, consequently, Terra, reflecting a broader market fear and uncertainty.

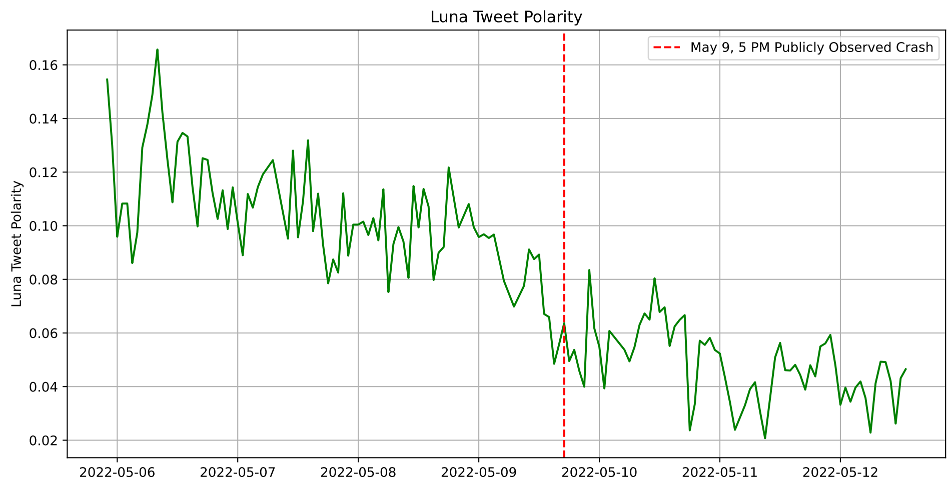

Figure 11: LUNA Tweet Sentiment : This figure shows the tweet sentiment for LUNA during the period: May 6, 2022, to May 12, 2022. A decreasing sentiment is observed for LUNA Post May 8. Tweets are collected from various social media platforms such as Reddit, Instagram, and Twitter.

Figure 12 captures investor disagreement about LUNA from May 6 to May 12, 2022. The standard deviation of tweet polarity highlights the extent of investor disagreement and volatility in opinions. Analyzing SD trends alongside average polarity scores offers deeper insights into overall sentiment direction and investor behavior. Polarity Mean is calculated to provide a weighted average polarity score for each hour, considering positive, negative, and neutral tweets about LUNA. Then, I compute the standard deviation of tweet polarities every hour to capture investor disagreement during that hour.

$$\text{Polarity\_Mean} = \frac{\text{Positive\_Tweets} - \text{Negative\_Tweets}}{\text{Positive\_Tweets} + \text{Negative\_Tweets} + \text{Neutral\_Tweets}}$$

Before May 9, the standard deviation is stable but slightly increasing (0.675 to 0.7), indicating moderate investor disagreement. There was a baseline level of uncertainty before the publicized crash. During the Terra crisis, the SD increased sharply and stayed elevated consistent with differences of opinion which is a precursor to adverse selection.

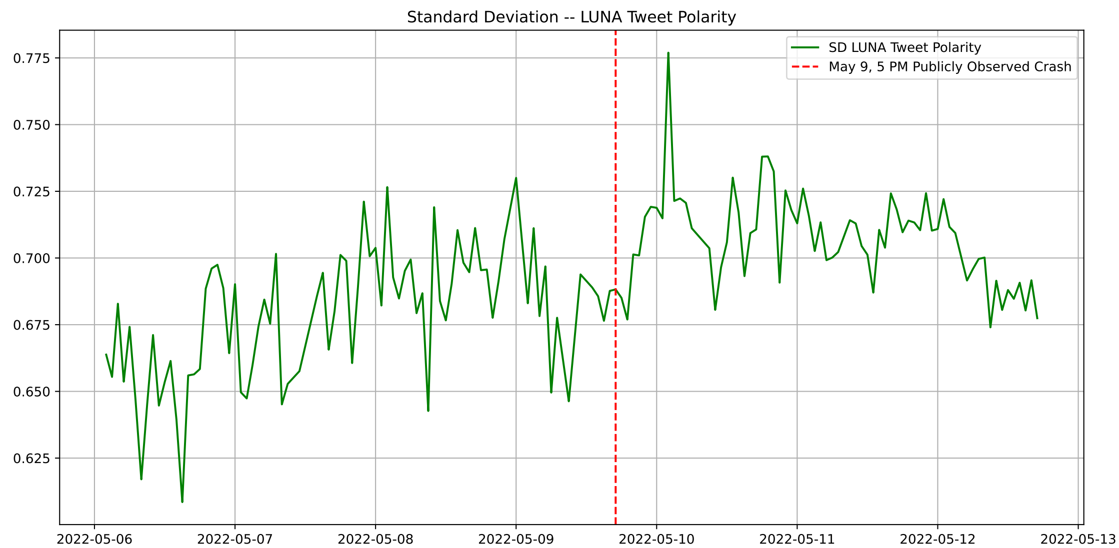

Figure 12: SD Tweet Polarity : This figure shows the standard deviation of Tweet Polarity for LUNA. This measure captures investor disagreement for LUNA during the period: May 6, 2022, to May 12, 2022. An increasing disagreement is observed for LUNA Post May 7. Tweets are collected from various social media platforms such as Reddit, Instagram, and Twitter.

Overall, the drastic decline in sentiment polarity and the sharp increase in SD after May 9 reflect a sudden and unexpected market reaction. The peak in SD around 5 PM on May 9, reaching approximately 0.75, aligns with the public acknowledgment of Terra's crash, showing that investor disagreement was only widely recognized at this point. Therefore, we cannot detect information diffusing into the market before May 9.

9 Conclusion

The Terra-LUNA crash wiped out approximately $50 billion, or 90% of its market value, within a week in May 2022. Our empirical analysis suggests that the collapse of Terra was driven by concentrated market manipulation by an investor or a group of investors that led to the breakdown of the arbitrage mechanism.

To investigate the cause of the crash, we begin by developing a novel mechanism that explains the run on Terra in May 2022. By precisely timestamping the major Terra withdrawals from various wallets, including Wallet A, Wallet B, and TFL, on the blockchain, we link these withdrawals—ranging from $10M to $375M—to the breakdown of the arbitrage mechanism.

Next, we investigate information-based models to determine whether the collapse was driven by information-based panic events. We compute the adverse selection components for various stablecoins and precisely timestamp the events, including withdrawals from Wallet A, Wallet B, and TFL, on the blockchain, that led to an increase in these components. Additionally, we calculate the net dollar outflows from Terra and the flight to safety into other stablecoins during the crisis period. We also analyze investor sentiment on social media during this time.

Our findings suggest that all information-based events—including increased investor disagreement, decreased tweet polarity, net dollar outflows from Terra, flight to safety, and heightened information asymmetry (as indicated by adverse selection measure)—occurred after the breakdown of the arbitrage mechanism at noon May 7. These findings support our claim that Terra's crash was driven by limits to arbitrage, triggered by a $175 million withdrawal.

Table 5: Significant Events Around Terra's Crash

Date	Event Description
Jan 19	Do Kwon announced the launch of the Luna Foundation Guard (LFG) a non-profit organization "to build reserves supporting the UST peg amid volatile market condition".
Feb 22	Singapore-based Luna Foundation Guard (LFG) raises $1 billion through the sale of LUNA tokens to buy bitcoin for UST's reserve system, with Jump Crypto and Three Arrows Capital being the lead investors.
March 23	• Do Kwon tweets "By my hand DAI will die" as he begins in earnest plans to starve off decentralized stablecoin DAI's liquidity on Curve. • Jump Trading, one of the investors behind LFG, proposes a mechanism for how to deploy the bitcoin (BTC) reserves to prop up UST's price in a crisis.

<table>
<tr><th>Date</th><th>Event Description</th></tr>
<tr><td>May 7</td><td>The first signs of the run reported in several industry reports and on social media.

Two addresses dubbed wallet A and B withdrew 400M UST from Anchor.
Wallet A was the first to withdraw 45M UST around 5am UTC and it sent the funds to Binance.
Following this event, wallet B withdrew 175M UST around noon and sent the funds to Ethereum using the Wormhole bridge.
Next, wallet A withdrew another 35M around 5pm, and then another 20M around 8:30pm and 9:30pm, again sending all the funds to Binance.
TFL removed 150M UST from the UST-3Crv pool at 21:44 PM to send it to a new UST-4Crv pool.
Finally, at 21:48pm wallet A withdrew the last 85M and send them to Curve.
</td></tr>
</table>

Date	Event Description
May 8	• LFG commits to loaning $750 million of BTC to market makers to defend the peg of UST and another $750 million of UST to be used to buy back BTC after volatility subsides. • Do Kwon jokes his way out of UST's depegging risk.
May 9	• Deposits on the Anchor protocol plunge below $9 billion from $14 billion after UST struggles to recover to $1. ANC, the protocol's token, fell 35% during the day. • UST loses its $1 peg for the second time and falls to as low as 35 cents. Do Kwon tweets @ 11:36 AM, "Deploying more capital – steady lads."
May 10	Claims that UST's depeg is due to a Soros-esque attack begin to emerge.
May 11	More than half, 58%, of traders place futures bets on higher LUNA prices despite Tuesday's drop, leading to $63 million in liquidations.
May 11	LUNA reaches price levels previously seen in August 2021. Value locked on Anchor, Terra's largest decentralized finance (DeFi) protocol, drops $11 billion over two days.

Date	Event Description
May 11	Do Kwon is revealed to be one of the pseudonymous co-founders behind the failed algorithmic stablecoin Basis Cash, CoinDesk reports.
May 12	The LUNA price falls 96% in a day, pushing it to less than 10 cents. The Terra blockchain is officially halted.

References

Angeletos, G.-M. and I. Werning (2006). Crises and prices: Information aggregation, multiplicity, and volatility. *american economic review 96*(5), 1720–1736.

Brancati, E. and M. Macchiavelli (2019). The information sensitivity of debt in good and bad times. *Journal of Financial Economics 133*(1), 99–112.

Brennan, M. J. and A. Subrahmanyam (1995). Investment analysis and price formation in securities markets. *Journal of financial economics 38*(3), 361–381.

Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 2007–2008. *Journal of Economic perspectives 23*(1), 77–100.

Catalini, C. and A. de Gortari (2021). On the economic design of stablecoins. *Available at SSRN 3899499*.

Chari, V. V. and R. Jagannathan (1988). Banking panics, information, and rational expectations equilibrium. *The Journal of Finance 43*(3), 749–761.

Cookson, J. A., C. Fox, J. Gil-Bazo, J. F. Imbet, and C. Schiller (2023). Social media as a bank run catalyst. *Available at SSRN 4422754*.

Dang, T. V., G. Gorton, and B. Holmstrom (2009). Opacity and the optimality of debt for liquidity provision. *Manuscript Yale University*, 1005–1029.

Dang, T. V., G. Gorton, and B. Holmström (2013). The information sensitivity of a security. *Unpublished working paper, Yale University*, 39–65.

Dang, T. V., G. Gorton, and B. Holmström (2020). The information view of financial crises. *Annual Review of Financial Economics 12*, 39–65.

Dang, T. V., G. Gorton, B. Holmström, and G. Ordonez (2017). Banks as secret keepers. *American Economic Review 107*(4), 1005–1029.

d'Avernas, A., V. Maurin, and Q. Vandeweyer (2022). Can stablecoins be stable? *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2022-131).

Diamond, D. W. (2023). Nobel lecture: Financial intermediaries and financial crises. *Journal of Political Economy 131*(10), 2597–2622.

Diamond, D. W. and P. H. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of political economy 91*(3), 401–419.

Easley, D., M. M. L. De Prado, and M. O'Hara (2011). The microstructure of the "flash crash": flow toxicity, liquidity crashes, and the probability of informed trading. *The Journal of Portfolio Management 37*(2), 118–128.

Easley, D., M. O'Hara, and L. Yang (2016). Differential access to price information in financial markets. *Journal of Financial and Quantitative Analysis 51*(4), 1071–1110.

Foley-Fisher, N., G. B. Gorton, and S. Verani (2020). Adverse selection dynamics in privately-produced safe debt markets.

Friedman, M. and A. J. Schwartz (2008). *A monetary history of the United States, 1867-1960*, Volume 9. Princeton University Press.

Gallagher, E. A., L. D. Schmidt, A. Timmermann, and R. Wermers (2020). Investor information acquisition and money market fund risk rebalancing during the 2011–2012 eurozone crisis. *The Review of Financial Studies 33*(4), 1445–1483.

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

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

Gorton, G. (2020). The regulation of private money. *Journal of Money, Credit and Banking 52*(S1), 21–42.

Gorton, G. and G. Ordonez (2014). Collateral crises. *American Economic Review 104*(2), 343–378.

Gorton, G. and G. Pennacchi (1990). Financial intermediaries and liquidity creation. *The Journal of Finance 45*(1), 49–71.

Gorton, G. B., E. C. Klee, C. P. Ross, S. Y. Ross, and A. P. Vardoulakis (2022). Leverage and stablecoin pegs. Technical report, National Bureau of Economic Research.

Gorton, G. B. and J. Y. Zhang (2023). Taming wildcat stablecoins. *U. Chi. L. Rev. 90*, 909.

Gromb, D. and D. Vayanos (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of financial Economics 66*(2-3), 361–407.

Iyer, R. and M. Puri (2012). Understanding bank runs: The importance of depositor-bank relationships and networks. *American Economic Review 102*(4), 1414–1445.

Liao, G. Y. and J. Caramichael (2022). Stablecoins: Growth potential and impact on banking.

Lin, J.-C., G. C. Sanger, and G. G. Booth (1995). Trade size and components of the bid-ask spread. *The Review of Financial Studies 8*(4), 1153–1183.

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

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

Ma, Y., Y. Zeng, and A. L. Zhang (2023). Stablecoin runs and the centralization of arbitrage. *Available at SSRN 4398546*.

Makarov, I. and A. Schoar (2022). Cryptocurrencies and decentralized finance (defi). *Brookings Papers on Economic Activity 2022*(1), 141–215.

Monnet, C. and D. R. Sanches (2015). Private money and banking regulation. *Journal of Money, Credit and Banking 47*(6), 1031–1062.

Morris, S. and H. S. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 587–597.

Narayanan, R. and M. E. Rhodes (2022). When safe became risky: the information sensitivity of subprime rmbs during the financial crisis of 2007-2009. *Available at SSRN 4040319*.

Schmidt, L., A. Timmermann, and R. Wermers (2016). Runs on money market mutual funds. *American Economic Review 106*(9), 2625–2657.

Shleifer, A. and R. W. Vishny (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.

Van Ness, B. F., R. A. Van Ness, and R. S. Warr (2001). How well do adverse selection components measure adverse selection? *Financial Management*, 77–98.