Learning from Terra-Luna: A Simulation-Based Study on Stabilizing Algorithmic Stablecoins

Federico Calandra[a], Francesco P. Rossi[a], Francesco Fabris[b], Marco Bernardo[a]

[a]*Dipartimento di Scienze Pure e Applicate, Università di Urbino, Italy*
[b]*Dipartimento di Matematica, Informatica e Geoscienze, Università di Trieste, Italy*

Abstract

This paper investigates the dynamics behind the catastrophic collapse of the Terra-Luna ecosystem in May 2022, where the UST stablecoin de-pegged and lost nearly all its value. Using a custom-built simulation environment, we reproduce the free-market interactions and protocol mechanisms that triggered the crash, offering a new perspective on the vulnerabilities of algorithmic stablecoins based on the dual-token seigniorage model. Building on this analysis, we propose four stabilization strategies: two purely algorithmic mechanisms, designed to prevent de-pegging and mitigate token hyperinflation, and two hybrid approaches, introducing partial collateralization through USDT and BTC reserves. Our simulations show that these strategies significantly reduce collapse events under extreme market conditions, with the best improvements ranging from 60.6% to 95.8%. Additionally, we gained better control over LUNA's circulating supply. Our simulation environment showcases how each proposal effectively mitigates systemic risks and enhances stability, thus providing valuable insights for designing future decentralized algorithmic stablecoins that can withstand market crises.

Keywords: Decentralized finance, Algorithmic stablecoins, Terra-Luna ecosystem, Dual-token seigniorage model

1. Introduction

The *decentralized finance* (*DeFi*) environment introduces a new paradigm in the finance sector. It enables unbanked users with smartphones and Internet access to engage in financial activities like money transfer, lending, borrowing, or speculating on synthetic stocks or commodities, on a 24/7/365 basis, without the need to involve banks or other financial institutions [1]. DeFi platforms utilize public and decentralized cryptocurrencies as the primary medium of exchange within their ecosystems.

The volatility of these digital assets has forced the introduction of *fiat-collateralized stablecoins* (CS). They are cryptocurrencies backed by a fiat currency or a commodity – typically USD, but also EUR, CHF, JPY, RMB, KRW, or gold. They dynamically maintain the peg with the corresponding fiat currency through suitable processes, along with a certification that, for each collateralized token, there exists a corresponding amount of value of the fiat currency, deposited in a bank, that can be redeemed.

Nowadays, we count dozens of fiat-backed stablecoins traded on the main cryptocurrency exchanges, such as USDT, USDC, TUSD, BUSD, and others. Their capitalization and success are increasing over time. They are used as a safe haven of stability when traders want to exit the volatility turmoil of the market and control the keys of the owned tokens. This avoids the use of a centralized exchange, frequently targeted by hacking and thefts, which is the only one able to swap a cryptocurrency for a fiat currency.

The spirit of the DeFi philosophy is that of creating a new paradigm for a decentralized financial system, totally detached from (central) banks, financial institutions, and fiat currencies. Unfortunately, fiat-collateralized stablecoins do not comply with this philosophy, since by using them the DeFi sector remains bound to the traditional financial market via the fiat currency used as collateral.

Another important category within the stablecoin ecosystem is *crypto-collateralized stablecoins*, which are backed by cryptocurrencies rather than fiat currencies. Due to the volatility of cryptocurrencies, these stablecoins require over-collateralization, meaning that users must deposit collateral exceeding the value of the issued stablecoins. Users are incentivized to buy the stablecoin when its price falls below the target value, realigning it with the peg. This approach sticks more closely to the decentralized philosophy of DeFi as it avoids reliance on fiat currencies, but introduces new challenges regarding efficient collateral management and resilience under stress conditions.

It is at this point that *algorithmic stablecoins* (AS) come into play. They are coins or tokens whose price is formally anchored to fiat currency solely by using an algorithmic protocol. More precisely, there are two different kinds of pure algorithmic stablecoin pegging mechanisms: *rebasing* and *seigniorage*.

In the rebasing model [2], the stablecoin total supply is not fixed and is modified adaptively on a regular basis, directly on the wallets of all users. The general idea is that when the price of the AS is above the parity with the adopted fiat currency – say the USD – it is necessary to increase the supply, since the demand is high and the token is too scarce. On the contrary, when the price is below the parity, it is necessary to decrease the

supply, since the demand is low and the token is too abundant. This implies that a user with, say, 1 000 tokens today, could find a greater amount in their wallet the day after, say 1 100, if AS price > \$1, without any action by the user. On the contrary, if AS price < \$1, then the wallet content could be smaller, say 900 tokens. The *Ampleforth* protocol (AMPL) [2] is an example of this model.

The seigniorage – or dual-token – model [3] typically has two tokens: the AS token and the *governance token* (GT). The AS token is the algorithmic stablecoin, while the GT token is used to absorb the volatility of the AS token. This means that when AS price < \$1, we can profitably burn 1 AS for \$1 worth of GT inside the protocol, which we can sell on the market, yielding the difference as a profit. This reduces the total supply of AS and stabilizes the price. On the contrary, when AS price > \$1, burning \$1 worth of GT can mint 1 AS, which we can sell on the market, yielding the difference as a profit. This increases the total supply of AS and stabilizes the price.

The *Terra-Luna* – or Terra – protocol [3], based on the AS TerraUSD (UST) and the GT LUNA, is of this kind. It has been, at the same time, the most successful and the worst example of how to build a decentralized AS. The most successful because it globally collected almost \$60B of capitalization and \$20B of *Total Value Locked* (TVL) – i.e., amount of USD locked in smart contracts of the DeFi's protocols – in no more than 15 months between January 2021 and May 2022, in an unprecedented rash of money induced by the *Anchor protocol* platform [4], which ensured 20% of *Annual Percentage Yield* (APY) for users who were lending UST. The worst example because more than 90% of the entire market cap was lost in 7 days between May 9 and May 15, 2022, as a consequence of a disastrous collapse induced by an irreversible depeg of UST, which crashed its value to almost zero [5].

The main objective of this paper is to study the dynamics of dual-token ASs and analyze the conditions that lead to their collapse. Specifically, our main objective is to lay the groundwork for tackling a crucial question: Is it possible to create a fully decentralized stablecoin that does not rely on fiat currency as collateral and makes efficient use of alternative forms of collateral?

This paper offers new insights into the challenges and potential solutions surrounding the design of an AS, taking inspiration from the infamous rise and fall of the Terra-Luna ecosystem. By revisiting one of the most ambitious attempts in the DeFi space – an experiment that soared to nearly \$60B in market capitalization only to unravel in a matter of days – we aim to capture both the promise and pitfalls of AS protocols. Our goal is to explore what went wrong and to propose viable improvements that could lead to more resilient AS systems.

This study builds upon and extends our previous work in [6]. We begin by revisiting the Terra-Luna protocol mechanics (Section 2) and continue by introducing the first key contribution of our paper: a custom simulation environment built in Matlab® (Section 3). This environment reproduces the dynamics of free-market interactions, allowing users to trade UST and LUNA while exploring various scenarios. By adjusting key parameters, we can simulate the conditions that precipitated UST's catastrophic collapse, offering a deeper understanding of the vulnerabilities inherent in seigniorage-based algorithmic stabilization mechanisms.

The second major contribution of our paper lies in the development and evaluation of four distinct stabilization proposals aimed at improving the resilience of AS protocols (Section 4).

Our first proposal introduces a re-engineered version of Terra's original stabilization algorithm, addressing some of the flaws that led to the collapse. Our simulations suggest that this redesign achieves a significant reduction in collapse events during extreme market conditions, demonstrating the potential for enhanced stability.

The second proposal introduces a purely algorithmic queue-based mechanism designed to prevent GT hyperinflation. This approach operates without external collateral, offering a novel path forward for AS protocols that seek to maintain stability while minimizing reliance on traditional financial instruments.

Next, we propose a third solution involving the establishment of a USDT reserve pool to support the AS token during market stress. This reserve only partially collateralizes the AS token, ensuring that the protocol retains its decentralized character while benefiting from added stability. We experiment with various collateralization levels to assess their impact on the overall system stability.

Fourthly, we address again GT hyperinflation – one of the core issues that plagued Terra – by implementing an automated bitcoin (BTC) reserve pool. Similar to the previous approach, it entails a partial collateralization of the AS. However, unlike the previous solution, it does not entirely rely on fiat-backed stablecoins, aligning more closely with the principles of DeFi.

We then present the outcomes of our simulations, illustrating the performance of each stabilization strategy under a variety of market conditions (Section 5) such as normal and panic scenarios. The paper concludes with a critical discussion of our findings, highlighting limitations, potential real-world implications, and avenues for future research to enhance AS stability and mitigate systemic risks (Section 6).

2. Background

In this section, we recall the basics of the Terra-Luna ecosystem. In Section 2.1, we introduce the concept of Automated Market Makers (AMMs). Then, in Section 2.2 we discuss the functioning of the Terra algorithmic market module and its role in maintaining the price stability of UST. Next, in Section 2.3 we analyze the key events that led to the collapse of Terra-Luna. Finally, in Section 2.4 we provide an overview of existing research on stablecoins and their significance for financial stability.

2.1. Automated Market Makers (AMMs)

AMMs, which serve as foundational components in *decentralized exchanges* (DEX) [7] by allowing the two tokens associated with them to be swapped, are revolutionizing assets trading through automated decentralized processes. Unlike traditional order-book-based exchanges, AMMs rely on *liquidity*

Table 1: Notation summary for the background section

Symbol/Term	Description
AMM	Automated Market Maker
DEX	Decentralized Exchange
LP	Liquidity Pool
VLP	Virtual Liquidity Pool
x, y	Reserve balances of the two tokens in a liquidity pool
k	Invariant in the constant-product formula ($k = x \cdot y$)
CP	Invariant for the VLP
$Pool_{Base}$	Initial quantity of UST in the VLP
$Price_{LUNA}$	Market price of LUNA (in USD), updated via oracles
δ	Deviation of the UST amount in the VLP from $Pool_{Base}$
$PoolRecoveryPeriod$	Inverse of the rate at which the VLP is replenished.
$TerraPool_{\delta}$	Stabilization mechanism accounting for UST supply deviation

pools (LP), constituted by the reserve of the two tokens, that algorithmically pair and maintain the two assets. A widely adopted pricing mechanism for AMMs is the *constant-product* formula [8], which ensures that any trade maintains the product of the token balances constant: $k = x \cdot y$, where x and y are the token reserve balances of the two tokens, while k is a constant called the *invariant* of the pool.

2.2. The Terra Stabilization Mechanism

The *Terra algorithmic market module* (TMM) played a central role in maintaining the price stability of UST. This is the module that provides incentives for arbitrageurs to mint or burn UST in response to price deviations from the peg [9]. An *arbitrageur* is an individual or entity that engages in the practice of exploiting price discrepancies in different markets to make profits.

When the UST's market price falls below the peg, e.g., \$0.98, arbitrageurs can profitably burn 1 UST obtaining automatically \$1 worth of LUNA from the protocol, making a \$0.02 profit per UST burnt. Conversely, if the UST's price exceeds the peg, e.g., \$1.02, they can burn \$1 worth of LUNA and mint 1 UST, again yielding a \$0.02 profit. The buying or selling pressure on UST generated by arbitrageurs helps realign the value of the stablecoin to the established level of one dollar. This realignment is further supported by the *principle of scarcity*: the minting or burning of UST affects its market availability. Specifically, when UST becomes scarcer, its value tends to increase, while excessive availability can lead to a decrease in value. A typical example of the application of this principle is represented by gold and Bitcoin, both of which maintain high value due to their rarity.

The mechanism just described operates via the protocol's algorithmic market-maker, the *virtual liquidity pool* (VLP), with LUNA's price sourced from validator oracles. Figure 1 illustrates the steps arbitrageurs take to profit during a de-peg event.

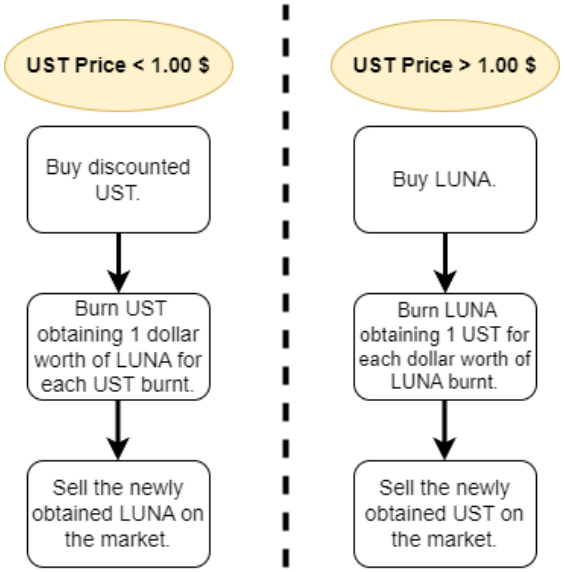

Figure 1: Strategy used by arbitrageurs in the Terra stabilization protocol.

The VLP is implemented by using a variant of the classical constant-product market-making algorithm [8]. In our formulation, the invariant CP of the pool is defined as:

$$CP = Pool_{Base}^2 \cdot \frac{1}{Price_{LUNA}} \tag{1}$$

$Pool_{Base}$ denotes the initial amount of UST allocated in the pool. The term $1/Price_{LUNA}$ adjusts the invariant in response to fluctuations in the external market price of LUNA (in USD), as reported by validator oracles. This term ensures that if $Price_{LUNA}$ increases, the effective value of the invariant decreases, and vice versa. This dynamic adjustment is critical as it allows the virtual liquidity pool to continually rebalance its asset composition in accordance with real-time market conditions.

The TMM integrates the $TerraPool_{\delta}$ stabilization mechanism, with the parameter δ indicating the deviation of the UST amount in the VLP compared to its base size $Pool_{Base}$:

$$Pool_{UST} = Pool_{Base} + \delta \qquad Pool_{LUNA} = \frac{CP}{Pool_{UST}} \tag{2}$$

The dynamics of δ play a crucial role in adjusting the liquidity pool sizes in response to market activities. As swaps happen and the balance between UST and LUNA quantities shifts, δ changes to reflect the deviation from the base pool size. A key aspect of the market module functionality is its ability to replenish the VLP, progressively bringing δ back towards zero. The rate of this replenishment is determined by the $PoolRecoveryPeriod$ parameter, defined in terms of blocks. At the end of each block – with one block being produced approximately every 6 seconds – δ is updated by changing it to:

$$\delta \cdot \left(1 - \frac{1}{PoolRecoveryPeriod}\right) \tag{3}$$

This formula governs the adjustment of δ, with $PoolRecoveryPeriod$ influencing the pace of the adjustment.

For example, let $Pool_{Base} = 1\,000\,000$ UST and assume $Price_{LUNA} = \$100$. Then:

$$CP = \frac{Pool_{Base}^2}{Price_{LUNA}} = \frac{(10^6)^2}{100} = 10^{10}$$

If heavy redemptions drive $\delta = 200\,000$, the pool balances become:

$$Pool_{UST} = 1\,200\,000 \qquad Pool_{LUNA} = \frac{CP}{Pool_{UST}} \approx 8\,333.33$$

A small swap of 1 UST then yields approximately:

$$\frac{Pool_{LUNA}}{Pool_{UST}} = \frac{8\,333.33}{1\,200\,000} \approx 0.00694 \text{ LUNA}$$

worth about \$0.694, instead of the ideal $\frac{1}{100} = 0.01$ LUNA (\$1). After each block, δ is multiplied by $(1-1/PoolRecoveryPeriod)$, gradually restoring $Pool_{UST}$ toward $Pool_{Base}$ and moving the swap yield back toward the full \$1 per UST.

The choice of *PoolRecoveryPeriod* thus involves a critical trade-off. A small *PoolRecoveryPeriod* (fast replenishment) means the pool can reset δ quickly, but this exposes the governance token to the risk of severe dilution and inflationary pressure. Conversely, a large *PoolRecoveryPeriod* (slow replenishment) limits sudden token inflation but may leave the peg vulnerable in periods of high volatility, as the VLP cannot restore its capacity fast enough to defend UST stability during a bank-run scenario.

PoolRecoveryPeriod was determined by the Terra community and, just before the time of the de-pegging event, its value was 36, meaning that a partial replenishment of the VLP occurred every $36 \cdot 6 = 216$ seconds if no transactions took place during this period [9]. Note, as a consequence, that a full replenishment can be obtained only when the number of blocks tends to infinity.

2.3. The Terra-Luna Collapse

The collapse of the Terra protocol was triggered by a complex series of events [10]:

1. On May 5, 2022, there was an evident selling pressure on UST and LUNA, indicated by negative hourly log returns. This selling pressure persisted, witnessing a loss of confidence in both cryptocurrencies.
2. On May 7, 2022, the stablecoin UST lost its peg with USD after a coordinated attack on Curve Finance's *Curve-3pool*. In this attack, Terraform Labs first transferred 150 million UST from the Curve-3CRV pool to the newly established Curve-4pool, which substantially reduced the liquidity available in the original pool. Shortly thereafter, an anonymous trader executed a swap of 85 million UST for USDC, exploiting the resulting imbalance. This sudden shift increased UST's share in the pool from approximately 33% to 66% and led to an immediate depreciation of UST to about \$0.985. This liquidity pool disruption set the stage for further destabilization, prompting an intervention from the *Luna Foundation Guard* (LFG) to temporarily defend the peg.
3. Despite this intervention, on May 9, 2022, UST lost its peg for the second and final time, leading to a significant decrease in both LUNA and UST prices. This event marked a critical blow to the stability of the Terra-Luna ecosystem.
4. Finally, on May 11, 2022, an announcement by Do Kwon (co-founder and CEO of *Terraform Labs*), presumably representing the last attempt to defend the peg by endorsing community proposal 1164, was interpreted by the market as a signal of the impending demise of the Terra-Luna ecosystem. This interpretation caused a final crash in the market, indicating the collapse of the Terra protocol.

These events highlight a cascade of failures within the Terra-Luna protocol, including an irreversible depeg of the UST stablecoin, the loss of over \$50B of capitalization [11], and the collapse of the other two DeFi platforms involved in the Terra-Luna ecosystem. They are the lending platform *Anchor* and the exchange *Mirror* protocol, which allowed users to create and trade *mirrored assets* (mAssets) that mirror the price of stocks and/or commodities of the real world.

The collapse of the Terra-Luna system can also be interpreted through the lens of a "bank run" phenomenon. In traditional banking, a crisis of confidence leads depositors to simultaneously withdraw their funds, depleting available liquidity and triggering a vicious cycle that accelerates institutional failure. Similarly, in the context of Terra-Luna, the sharp drop in LUNA's market capitalization – which served as the assets backing the liabilities represented by the circulating UST – led to a significant loss of market confidence. As users perceived that the value of the backing assets (LUNA) was rapidly declining relative to the outstanding obligations (UST), they responded by massively selling their holdings. This exacerbated the depeg and fueled the self-reinforcing downward spiral typical of a bank run.

2.4. Literature Review

Stablecoins have gained interest from researchers across various disciplines as well as from financial institutions. The literature delves into their design, mechanisms, and impacts on financial stability, monetary policy, and regulation.

The European Central Bank (ECB) *Crypto-Assets Task Force* tackles the issue of stablecoins in its report n. 247 [12]. Calcaterra et al. [13] study the first-order design principles for stablecoins by illustrating the core design features and their interoperative feedback. The recent Bank for International Settlements (BIS) paper n. 141 [14] provides an overview of the evolution of the stablecoin market over the past decade and examines whether stablecoins have stayed true to their name in terms of being "stable". Ante et al. [15] analyze 22 peer-reviewed articles, offering insights into stablecoin types, benefits, risks, and regulatory challenges. They identify research gaps, notably the lack of robust data and frameworks for analyzing the stablecoin ecosystem complexities. Furthermore, recent research by Choi and Kim [16] examines the challenges and opportunities associated with stablecoins and *central bank digital currency* (CBDC). They explore the financial stability implications of stablecoins and the potential for CBDCs to revolutionize payment infrastructures and monetary policies. As for ASs, Clements [17] examines their fragility, highlighting risks like self-fulfilling runs and coordination failures. Recent collapses, such as the Terra-Luna incident, bolster Clements' argument, prompting questions about designing stablecoins resilient to high volatility.

The specific ecosystem of Terra-Luna has been the subject of several studies that explore its unique features and challenges, as well as the failure that occurred in May 2022.

Cho [18] explains the collapse of the Terra project by analyzing the impact of the Anchor protocol on the system stability. Anchor was a great catalyzer for the demand and supply of UST, but – at the same time – was one of the responsible for the collapse of the system. Cho shows that, during the de-pegging event, the UST supply rapidly increased from 2.3B to 3.4B, contrary to the expected contraction mechanism that should have reduced the UST supply when the UST price was below the peg. The article attributes this anomaly to the run on the Anchor protocol, which allowed users to borrow UST at a low interest rate and sell it on the market, creating a downward pressure on the UST price and an upward pressure on the UST supply. The article of Cho shows that the UST price on exchanges followed the redeemed value of UST that users could obtain by swapping UST for LUNA and selling it on the market. Cho finds that the redeemed value of UST was consistently lower than the UST price on exchanges, indicating that users were undercompensated when they redeemed UST for LUNA. Cho also finds that the UST price on exchanges followed the redeemed value of UST closely, suggesting that users were arbitraging the price difference by selling UST on the market and buying UST on the Terra blockchain. This mechanism played a significant role in the de-pegging event, as it created a disincentive for users to hold UST and a downward pressure on the UST price.

Briola et al. [10] quantitatively describe the main events that led to the Terra project failure by reviewing, in a systematic way, news from heterogeneous social media sources and discussing the fragility of the Terra project and its vicious dependence on the Anchor protocol. They also identify the crash trigger events by analyzing hourly and transaction data for BTC, LUNA, and UST.

Liu et al. [11] use data from the Terra blockchain and trading data from exchanges to study the dynamics and interactions of the system. They show that it was a complex phenomenon that happened across multiple chains and assets and that the run on Terra was not due to market manipulation, but rather to growing concerns about the sustainability of the system.

Kurovskiy and Rostova [19] investigate the collapse of the Terra-Luna ecosystem by using transaction-level data from the Terra blockchain and cryptocurrency exchanges, illustrating the several flaws in the design of UST that impeded its price stabilization.

Uhlig [5] develops a novel theory to account for the Terra crash and uses it to shed light on data. He presents a new methodology to show how crashes unfold gradually, by introducing the method of quantitative interpretation.

Ferretti and Furini [20] investigate the collapse of UST through Twitter as a passive sensor, by analyzing sentiment in tweets to explore correlations with market value and highlighting the challenge of foreseeing sudden catastrophic events solely through sentiment analysis.

Recent work on bank runs also offers valuable insights into the fragility of algorithmic stablecoins. Adams and Ibert [21] show how flaws in arbitrage mechanisms can trigger rapid runs, as seen in the collapse of Iron Finance – an algorithmic stablecoin protocol launched on the Polygon blockchain in 2021. Saengchote and Samphantharak [22] draw parallels between Iron Finance and TerraUSD, highlighting confidence erosion as a key driver of destabilization. Krogstrup et al. [23] emphasize the role of digital technologies in accelerating run dynamics.

3. The Simulation Environment

Now we present two distinct and independent simulations of the Terra stability mechanism. The first one investigates the role of the mechanism in the collapse of Terra-Luna and is rooted in Cho's study [18], which highlights the system's limited redemption capacity. The second model delves into the repercussions of a sudden and substantial increase in the GT token supply during a collapse, resulting in its complete devaluation over a short period.

The simulation environment is designed primarily to investigate the effectiveness of different stability mechanisms for ASs and compare them. To achieve this, the model integrates and parametrizes several exogenous elements – such as token price dynamics, liquidity conditions, and stochastic user behaviors – that are intended to mimic real-market scenarios. It is important to note that these parameters have been chosen in an ad hoc manner, based on theoretical considerations and our preliminary understanding of market mechanisms. Fine-tuning these parameters with real-world data will be the focus of future studies, which will aim to validate and enhance the robustness of the model.

The simulation environment will be used in Section 4 to propose four enhancements to the Terra protocol. The first one consists of a redesign of the stability mechanism, while the other three have to do with limiting the LUNA supply growth.

In Section 3.1 we discuss the implementation of a mechanism to model the price dynamics of the AS and GT tokens. In Section 3.2 we explain how we estimate the distribution of tokens among users' wallets, drawing parallels from existing cryptocurrencies to inform our simulation. In Section 3.3 we introduce the probabilistic approach of token swaps within liquidity pools, utilizing a random walk model to account for market volatility. In Section 3.4 we describe the implementation of the arbitrage mechanism within our simulation. Finally, in Section 3.5 we explore the mechanisms underlying panic selling and irrational buying behaviors in the market, defining specific zones that trigger these behaviors during simulated market fluctuations.

3.1. Price Dynamics through an AMM

To achieve our goals effectively, we need an environment that can accurately reflect the price changes in the free markets of the exchanges, for both the AS and the GT tokens, and the dynamics of the Terra stabilizing protocol. This can be done by simulating an AMM, which operates within discrete time intervals called *iterations* or *samples*.

In our simulations, two AMMs (or markets) are implemented: one that governs the buying and selling of the AS token and the other that governs the buying and selling of the GT token. During each time interval, a swap occurs within each

Table 2: Notation summary for the simulation environment

Symbol/Term	Description
Π_{T_a,T_b}	Liquidity pool containing tokens T_a and T_b
T_a, T_b	Tokens traded in the pool (e.g., T_s, T_v, or T_U)
$Q_a(n)$, $Q_b(n)$	Respective token balances in the pool at iteration n
k	Invariant of the pool: $k = Q_a(n) \cdot Q_b(n)$
q_a, q_b	Input and output token quantities in a swap
$P_a(n)$	Price of token T_a at iteration n, given by $P_a(n) = \frac{Q_U(n)}{Q_a(n)}$
T_U	Fully collateralized stablecoin pegged to USD
T_s	Algorithmic stablecoin token
T_v	Volatile token (e.g., LUNA)
$p(n)$	Sell probability at iteration n in the stochastic swap model
Δ	Stochastic variable representing change in sell probability
μ_Δ, σ_Δ	Mean and standard deviation of Δ, with $\mu_\Delta = 0$ at equilibrium
ρ	Tolerance threshold for healthy zone deviations in T_s price
lowerBound, *upperBound*	Bounds defining the healthy zone for the sell probability $p(n)$
ξ	Deviation from parity used in calculating $P_{Terra}(\xi)$
a, t	Constants in the definition of the Terra protocol usage probability

AMM, influencing simulated token prices. The impact of these swaps on prices depends on the volume of involved tokens: larger volumes have a greater effect on simulated prices, thus inducing *slippage* – a difference between the expected price of a trade and the price at which the trade is actually executed, typically caused by limited liquidity or large trade sizes relative to the pool.

Let us explore in detail how AMMs are implemented within the simulations. Each AMM is described by an LP Π_{T_a,T_b} composed of two tokens T_a and T_b. At the discrete time instant n, the state of the LP is defined by:

- $Q_a(n)$ and $Q_b(n)$, which represent the token balances in the pools of T_a and T_b, respectively, at iteration n.
- $k = Q_a(n) \cdot Q_b(n)$, which is the invariant at time n.

In our simulations, we assume zero transaction fees and constant pool liquidity for simplicity, as these factors are deemed unimportant for our analysis goals. The LP state at a given time n can be written as:

$$\Pi_{T_a,T_b}(Q_a(n), Q_b(n))$$

We can define a swap as a function that operates on the LP state. At time n, a swap occurs when a trader sells q_a units of token T_a into the pool and buys q_b units of token T_b from the pool. Concretely, after the swap, the new reserves are:

$$Q_a(n+1) = Q_a(n) + q_a \qquad Q_b(n+1) = Q_b(n) - q_b$$

Solving for the output amount q_b at time $n+1$ yields:

$$q_b \;=\; Q_b(n) - Q_b(n+1) \;=\; \frac{k}{Q_a(n)} \;-\; \frac{k}{Q_a(n)+q_a} \tag{4}$$

where $k = Q_a(n) \cdot Q_b(n)$ is the constant-product invariant.

In an AMM, the token price is measured in terms of the other token in the LP. Therefore, it is crucial to establish a fixed reference for token pricing. Here, the reference is USD, serving as a stable benchmark for analyzing price fluctuations of the algorithmic stablecoin. As fiat currency cannot be used in DeFi, we need to make an assumption. Let us designate the other token as T_U, a fully collateralized stablecoin pegged to USD; for example, USDT or USDC. This allows token values to be expressed in USD, assuming that T_U maintains a constant external value of \$1 throughout the simulation.

The concept of AS involves a stabilization mechanism that is usually based on two tokens: the AS token (e.g., UST) – which will be referred to as T_s – and the GT token, which is volatile (e.g., LUNA) – from now on T_v. Both simulations instantiate the AMM model above by employing two LPs to emulate the price of the T_s and T_v tokens. The first pool, denoted by Π^S, is designed to simulate the market operations on T_s and consists of T_s and T_U. The second pool, denoted by Π^V, replicates the market dynamics of T_v and consists of T_v and T_U.

Although real-world stablecoin-to-stablecoin pools (e.g., Curve) use tailored, low-slippage designs, we employ a standard constant-product AMM for Π^S to model UST price dynamics. This simple mechanism captures key price-impact effects with minimal calibration; investigating more complex pool architectures is left for future work.

At each time instant n, a swap occurs within Π^S and Π^V thus altering the prices of T_s and T_v, the state of the LP is given by $\Pi_{T_a,T_U}(Q_a(n), Q_U(n))$ where $T_a = T_s$ or $T_a = T_v$, and the price $P_a(n)$ of the token T_a is determined by the ratio of the quantities of the two tokens inside the pool at that specific moment:

$$P_a(n) = \frac{Q_U(n)}{Q_a(n)} \tag{5}$$

This is what Uniswap, the most used DEX with the highest TVL, refers to as the *mid price* [24]. It can be viewed as the price at which one could theoretically trade an infinitesimally small amount of one token for the other in the pool, without slippage of the price.

3.2. Wallet Distribution

The simulation environment imposes the crucial choice of setting the effective quantities of each token we have to swap to simulate an ordinary session of the free market or the Terra-Luna redemption protocol. One possible approach could be that of using information about the amount of tokens inside the wallet of each user. Since there are no data about the original distribution of UST and LUNA, we estimated such data from the distribution of BTC and ETH among addresses within their respective blockchains [25], under the hypothesis that all the cryptocurrencies show a similar wallet balance distribution. We found that BTC and ETH balances can be approximated with an exponential distribution with parameter λ, whose probability density function is defined as follows:

$$f(B,\lambda) = \begin{cases} \lambda e^{-\lambda B} & \text{if } B > 0 \\ 0 & \text{if } B \le 0 \end{cases}$$

where B represents a given wallet balance and parameter λ was obtained using MATLAB's `fitdist()` function, which returns the value of λ that best fits the provided input data – in our case, the balances of BTC wallets. The analysis reveals that the majority of the wealth is held by a few "whales". For example, at the time of writing over 90% of all Bitcoins are held in wallets with a balance of more than 1 BTC, with only 4 addresses holding more than 100 000 BTC [26].

In each iteration of the simulation, a wallet is randomly selected and its available balance B is determined. The trade size q is then sampled from a zero-mean normal distribution, where the standard deviation is a fraction of B (adjusted by using a volatility scaling parameter). This distribution is truncated to the interval $[0, B]$. Drawing q in this way ensures it is always non-negative and within the wallet balance, while the spread can widen during crisis scenarios to simulate panic trading. This approach allows the model to realistically reflect the variability and heterogeneity of swap sizes.

3.3. Stochastic Swaps

The probability of buying or selling a token within an LP is set by using the stochastic process of a *random walk* [27], which fully reflects the intrinsic uncertainty of the market. A random walk represents the cumulative effect of a sequence of random steps or movements, taken at discrete time intervals, starting from an initial position. Mathematically, a simple one-dimensional random walk can be defined as follows. Let $(X_i)_{i\in\mathbb{N}}$ be a sequence of independent and identically distributed random variables, with a probability distribution function $\mathcal{P}(X_i) = \mathcal{P}(X_0)$, representing the successive steps taken at discrete time points $t = 0, 1, 2, ...$ on the left or the right. The position of the walker at time t is given by the sum of all the random steps up to that point:

$$S_t = X_0 + X_1 + X_2 + \ldots + X_t$$

In the stochastic swap model presented in this study, a random walk is employed to determine the probability of a token T to be swapped within the pool at a given iteration. Let us consider the example of LP Π_{T_a,T_b}. At iteration n, $T = T_a$ with probability $p(n)$ and $T = T_b$ with probability $1 - p(n)$. So, there are three possible scenarios:

$$\begin{aligned} p(n) = \tfrac{1}{2} \quad &\Longrightarrow \quad \text{equilibrium condition} \\ p(n) < \tfrac{1}{2} \quad &\Longrightarrow \quad T_a \text{ experiences buying pressure} \\ p(n) > \tfrac{1}{2} \quad &\Longrightarrow \quad T_a \text{ experiences selling pressure} \end{aligned} \tag{6}$$

At the beginning of each iteration, $p(n)$ is varied from its current value by an amount Δ:

$$p(n+1) = p(n) + \Delta \tag{7}$$

where Δ is a random variable with a normal distribution $\Phi(\mu_\Delta, \sigma_\Delta^2)$ such that its mean value μ_Δ is equal to zero, while its variance σ_Δ^2 is a simulation parameter initialized during the simulation setup that expresses the market volatility. To ensure that the swap probability remains within the valid range, we truncate any value of $p(n+1)$ outside $[0, 1]$ to the nearest boundary.

At the beginning of the simulations, we suppose that the market is in a state of equilibrium, hence $p(0) = \frac{1}{2}$. The value of $p(n)$ is subsequently updated based on equation (7), allowing for both positive or negative steps, which increases or reduces the probability of sale. These price dynamics are illustrated in Figure 2, which shows the correlation between the token price fluctuations and the evolving sell probability. Each iteration proceeds as follows:

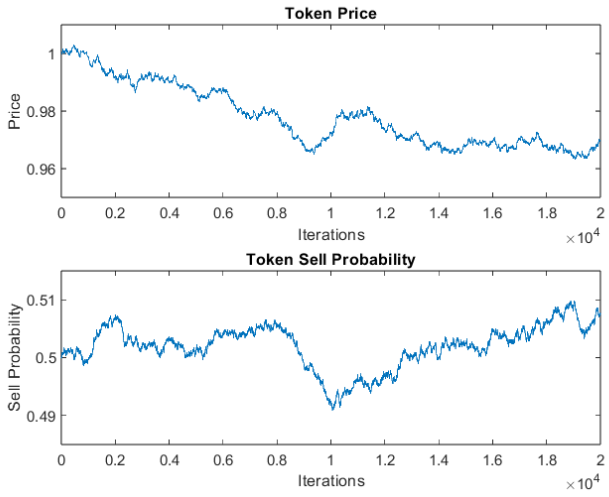

Figure 2: The top plot illustrates the token price fluctuations over time in a stochastic swap model based on a random walk. The bottom plot depicts the corresponding sell probability throughout the same period.

1. **Wallet selection**: choose one wallet at random.
2. **Swap size**: sample q from a Gaussian distribution truncated to $[0, B]$, where B is that wallet balance.
3. **Direction**: decide buy vs. sell for token T_a based on probability $p(n)$.
4. **Execution**: apply the swap of size q in the pool, updating Q_a and Q_b.

Although the stochastic swap processes for T_s and T_v evolve independently under normal market conditions, the Terra stabilization mechanism couples their price trajectories. Specifically, deviations in the price of T_s trigger minting or burning of T_v via arbitrage operations (Section 2.2). Thus, there is a dynamic correlation between the prices of T_s and T_v.

3.4. Arbitrage Mechanism Implementation

In each iteration, we simulate arbitrage by selecting a random wallet and estimating the profit from minting or burning based on the current price deviation of the stablecoin T_s. The arbitrageur uses the wallet's available T_U balance to execute the maximum profitable swap – burning T_s when its price exceeds the peg or burning the governance token T_v when the price falls below. This action updates both liquidity pools and embeds the protocol incentive mechanism directly into our simulation.

3.5. Inducing the Collapse

We now have the problem of representing the conditions of a FUD – *Fear, Uncertainty, Doubt* – panic sell or of an irrational FOMO – *Fear of Missing Out* – buy that can afflict the free market. During the FUD phase, users who are in a state of panic tend to sell their cryptocurrencies with great intensity, triggering a bank run. Widespread panic can be triggered by any external event. In the case of the Terra-Luna ecosystem, it was the withdrawal of large amounts of UST from the Anchor protocol and the partial loss of the peg.

We can model this by introducing the concept of *panic zone* for the stablecoin. When the market enters the panic zone, a mechanism is triggered to represent the irrational behavior of users, whose decisions are driven more by emotions than rationality. Otherwise, we are in the *healthy zone*, where the users act as usual, in a normal condition of the market. It is understood that the stablecoin tends to maintain the peg inside the healthy zone, while it tends to lose it outside, when entering the panic zone.

Concerning the pool Π^S, we have implemented two different definitions of panic zones, used in the simulations to represent the behavior of the buying/selling probability, when hovering over the boundaries separating the healthy zone from the panic zone of the market. The two approaches respectively use the price of the stablecoin (Section 3.5.1) and its probability of sale (Section 3.5.2) to trigger the risk of a collapse when entering the panic zone.

3.5.1. The Variable Mean Approach Based on Price

Suppose that at iteration $n = 0$ the price is set to $P_s(0) = 1$. As the simulation progresses, various buy and sell transactions of T_s will occur, leading to fluctuations in its probability and price based on equations (6) and (7). At iteration n, T_s is in the healthy zone if:

$$|1 - P_s(n)| < \rho \qquad \text{with } \rho < 1 \tag{8}$$

In our simulation we have set $\rho = 0.05$; this implies that the healthy zone of T_s covers the interval $\$0.95 < P_s(n) < \$1.05$. When the price of the algorithmic stablecoin falls below the lower bound of the healthy zone ($P_s(n) < \$0.95$), the system enters a panic zone, and strong selling pressure is triggered on T_s. In this case, we set

$$\mu_\Delta = (1 - P_s(n)) \cdot \sigma_\Delta,$$

which is positive and scales with the severity of the depeg, driving exponential sell-off. For all the other values ($P_s(n) \geq \$0.95$, including $P_s(n) > \$1.05$), we keep $\mu_\Delta = 0$, indicating no panic-driven pressure.

So, when $1 - P_s(n) \geq \rho$ we are in the panic zone and the following market condition holds:

$$1 - P_s(n) > 0 \implies \mu_\Delta > 0 \implies \text{selling pressure}$$

Our choice for $\rho = 0.05$ was arbitrary but conservative. Intuitively, narrowing this interval (e.g., to [0.98, 1.02]) would trigger the panic-driven mechanisms more frequently, potentially leading to a slight increase in the number of collapses. Conversely, widening the interval (e.g., to [0.90, 1.10]) would delay or reduce panic responses, lowering collapse frequency. However, empirical evidence supports the appropriateness of our choice: between January 1 and April 30, 2022 – the months preceding the collapse – 95% of UST's intraday prices fell within the interval [0.9913, 1.0114], based on historical prices [28] as illustrated in Figure 3.. Over the 120-day period, the lowest price recorded was \$0.9754 and the highest was \$1.0150. The standard deviation of daily lows and highs was only \$0.0031 and \$0.0025, respectively. This confirms that UST was highly stable during this time, and that our selected "healthy" range of [0.95, 1.05] provides a broad but realistic tolerance band without distorting comparative results across scenarios.

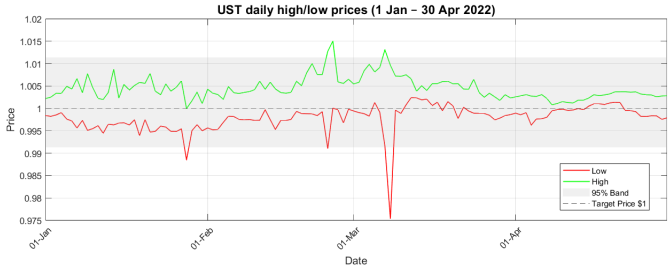

Figure 3: Daily high (green) and low (red) intraday prices of UST from January 1 to April 30, 2022. The shaded gray area marks the 95% band defined by the 2.5^{th} and 97.5^{th} percentiles of all daily lows and highs, and the dashed line indicates the \$1 peg. Throughout this period, 95% of intraday prices remained within [\$0.9913, \$1.0114].

These dynamics are also applied to the volatile token T_v, as seen in the May 2022 Terra ecosystem collapse, where LUNA was subjected to severe selling pressure after the UST depeg. In our implementation, whenever the stablecoin enters the panic zone, holders of T_v anticipate dilution – since the stabilization mechanism mints new T_v to burn T_s – and respond by selling T_v in proportion to the magnitude of the depeg.

This simulation captures the critical moment when the stablecoin exits the healthy zone and enters a state of collapse, illustrating how small deviations from the peg can lead to exponential selling pressure in both the stablecoin and its collateral token.

3.5.2. The Variable Probability Approach

In the second approach, the sell probability $p(n)$ freely fluctuates as described in equations (6) and (7), when we are in a healthy zone. The only difference is that now the healthy zone itself is defined in terms of probabilities, more precisely as the interval $lowerBound \leq p(n) \leq upperBound$. Outside this interval, we are in the panic zone and the user behavior is modeled through a completely different function representing their irrational mindset, which is an exponential-type function. In our case $lowerBound = 0.3$ and $upperBound = 0.7$, although these values can be adjusted as needed.

Define now the "$\mathcal{B}$ zone", or *buy zone*, as the market condition in which $p(n) < lowerBound$: in this scenario, the demand for buying cryptocurrency is particularly high for various reasons. On the contrary, the "$\mathcal{S}$ zone", or *sell zone*, represents the most critical phase of the market, in which $p(n) > upperBound$ and the stablecoin risks a collapse. The scheme illustrated in Figure 1 is typically not followed when the market enters zone $\mathcal{S}$. Under normal conditions, when the price of UST falls below \$1, this would present an opportunity to purchase discounted UST and profit from the arbitrage. However, when the drop is sharp and the market enters a panic phase, users are primarily focused on offloading their existing UST holdings rather than acquiring more to seek profit – particularly because the value of LUNA received in exchange is also rapidly declining.

The functions $\mathcal{B}(d)$ and $\mathcal{S}(d)$ compute the sell probability within the panic zone:

$$\mathcal{B}(d) = \frac{1}{1+e^{\frac{d+a_{\mathcal{B}}}{\eta}}} \qquad \mathcal{S}(d) = \frac{1}{1+e^{-\frac{d+a_{\mathcal{S}}}{\eta}}} \tag{9}$$

Here d is the distance from the healthy zone, η is a parameter determining the slope of the function, and $a_{\mathcal{B}}$ and $a_{\mathcal{S}}$ need to be computed based on the parameters *lowerBound* and *upperBound* fixing the limits of the healthy zone, so as to align the values of d inside and outside the interval. More precisely, we have:

$$a_{\mathcal{B}} = d \quad \text{if} \quad \frac{1}{1+e^{\frac{d}{\eta}}} = lowerBound$$

$$a_{\mathcal{S}} = d \quad \text{if} \quad \frac{1}{1+e^{-\frac{d}{\eta}}} = upperBound$$

In our case $a_{\mathcal{B}} = a_{\mathcal{S}}$ as $0.5 - lowerBound = upperBound - 0.5$.

Inside the Terra-Luna protocol, the simulation keeps track of the number of tokens in circulation for both cryptocurrencies and incorporates the principle of scarcity by modifying the evolution of the swap probability $p(n)$ within the market stochastic process. Whenever a token is minted, a positive deterministic shift – proportional to the minted quantity – is applied to $p(n)$, increasing the likelihood of a sale. Conversely, when a token is burnt, a negative deterministic shift is applied, increasing the likelihood of a purchase. These shifts are integrated into the underlying random walk that drives the market behavior, thereby directly linking token supply dynamics to changes in swap propensity.

The Terra-Luna protocol is not used during each time unit, but its probability of use $P_{Terra}(\xi)$ increases with the deviation ξ from the peg value. This is because a larger variation in the price of UST offers greater profit opportunities for arbitrageurs, who tend to use the protocol more frequently. The function $P_{Terra}(\xi)$ is defined as follows:

$$P_{Terra}(\xi) = e^{-\frac{1}{(\xi+a)^t}} \cdot 0.5 + 0.5 \tag{10}$$

with a and t suitable constants. In our case $a = 0.55$ and $t = 5$, while $\xi = |1 - \text{UST}.price| \cdot 10$ is the deviation from parity. Equation (10) assumes different interpretations depending on whether the price of UST is above or below the parity. In the former case, the function returns the probability of minting UST; in the latter case, the function gives the probability of burning UST.

4. Four Terra-Luna Stabilization Proposals

We now introduce four proposals aimed at improving the stability of the Terra-Luna protocol, each addressing different aspects of the system. In Section 4.1 we propose a modification to the VLP replenishing mechanism to make it more responsive during periods of peg loss. In Section 4.2 we present a purely algorithmic queue system to limit the excessive minting of LUNA without relying on external reserves. In Section 4.3 we introduce a USDT reserve pool to stabilize the price during crises. Finally, in Section 4.4 we describe an automated BTC reserve pool to counteract the hyperinflation of LUNA.

4.1. Modifying the VLP Replenishing Mechanism

At the core of Terra's stabilization mechanism is the VLP [9], which in our case corresponds to $\Pi^{virtual}_{T_s,T_v}$, i.e., the pool containing both the stable (UST) and the volatile (LUNA) tokens, respectively. Initially, it comprises an equal quantity of $Pool_{Base}$ units for T_s and T_v. In equation (3) we introduced the parameter δ, which is equal to the difference between the current quantity of T_s and the initial baseline quantity $Pool_{Base}$ of T_s in the VLP. Under these hypotheses, at each iteration n the invariant of the VLP, denoted by $CP(n)$, is defined as:

$$CP(n) \;=\; Pool^2_{Base} \;\cdot\; \frac{1}{P_v(n)}$$

where $P_v(n)$ is the market price of the volatile token T_v at time n. Since $P_v(n)$ evolves over time, the pool invariant $CP(n)$ also depends on n. The corresponding token reserves are then:

$$\begin{cases} Q_s(n) = Pool_{Base} + \delta(n) \\ \\ Q_v(n) = \frac{CP(n)}{Q_s(n)} \end{cases}$$

Each swap executed within the pool dynamically alters $\delta(n)$ and hence the reserves $Q_s(n)$, $Q_v(n)$.

In the original implementation of the pool replenishing mechanism of equation (3), $\delta(n)$ goes down to zero only when n tends to infinity. So, we suggest a different method that acts similarly, but can bring $\delta(n)$ to zero after exactly *PoolRecoveryPeriod* (*RP* for short) blocks. The idea is very simple. Let us consider the scenario where the initial swap within the pool at time $n = 0$ involves an amount of x stable tokens T_s. Then the replenishment of x tokens in the VLP is evenly spread among a list of *RP* chunks, which will determine δ in the subsequent *RP* time instants of the simulation. More formally, let A be an initially empty vector of length *RP*. At the time instant $n = 0$, after the initial swap of x T_s, it is filled with the *RP* chunks:

$$A(0) = \left[\left(\frac{x}{RP}\right), \left(\frac{x}{RP}\right), ..., \left(\frac{x}{RP}\right)\right] \qquad |A| = RP$$

as a consequence of the swap. The first element A_1 of A will determine the value of $\delta(1)$ as follows:

$$\delta(1) = \delta(0) - A_1$$

In general, when a swap operation takes place within the virtual pool – denoted by $Swap(T_i, q)$, where $T_i \in \{T_s, T_v\}$ – a quantity vector Q is initialized as $Q = [q/RP, q/RP, ..., q/RP]$, where $|Q| = RP$. Consequently, the update of vector A takes place according to the following procedure:

$$\begin{cases} A(n) = A(n-1) + Q & \text{if } T_i = T_s \\ A(n) = A(n-1) - Q & \text{if } T_i = T_v \end{cases} \tag{11}$$

After this, $\delta(n)$ is updated as follows, where now $A_1(n)$ is the first element of $A(n)$:

$$\delta(n) = \delta(n-1) - A_1(n)$$

Finally, the elements of the vector A are shifted left and the last element is set to zero:

$$A(n) = [A_1, A_2, ..., A_{RP}] \xrightarrow{shift} A(n+1) = [A_2, A_3, ..., A_{RP}, 0]$$

This process guarantees that, if no swaps occur, the VLP is fully replenished after exactly RP blocks. Based on this setup, the improvement of the stability we are going to propose pertains to a variant of the VLP replenishing mechanism suggested above.

If the price of T_s falls below the peg, we propose to reduce the length of the replenishing vector A dynamically. The algorithm modifies RP to accelerate the restoration of the VLP when the price significantly deviates from the peg. The replenishment period is adjusted by examining the price of the stable token (T_s) relative to a series of thresholds. The code checks these thresholds iteratively and, based on how far the price has dropped below the peg, it determines a new length for the restored vector A. As the price of T_s drops, RP is decreased, thereby speeding up the pool replenishment process (i.e., increasing the rate at which δ is reduced). The shortening of RP is controlled by a set of predefined price thresholds: in our simulation, we chose to range these thresholds from 0.950 to 0.995. For each crossed threshold, RP is reduced proportionally by shrinking the length of vector A. The procedure to dynamically adjust RP based on the stablecoin price can be described as follows:

1. Check the price: At each time step, the stablecoin price is compared against a series of predefined threshold values.
2. Adjust RP: Based on the extent to which the price has fallen below the peg, the algorithm reduces RP by a proportion of the initial pool recovery period.
3. Shrink restore vector: The length of A is modified by equally redistributing the remaining value of the excess restore chunks among the shortened vector. This ensures that the sum of the restore values remains consistent while accelerating the rate of restoration.

Essentially, in a crisis scenario, the VLP redemption capacity dynamically increases. Given that the variation of δ correlates directly with the price $P_s(n)$ of T_s (as shown in [18]), this modification enables the system to recover more rapidly from a peg loss.

4.2. Implementing a Stochastic Queue System against LUNA Hyperinflation

This approach aims to limit the excessive minting of LUNA through a purely algorithmic queue system that does not rely on any collateral. The hyperinflation of LUNA was one of the main causes of the system collapse. This queue system can be easily derived by referring to Figure 1: queue A gathers all transactions from users who wish to burn their UST in exchange for LUNA, while queue B collects transactions from those who want to burn their LUNA to obtain UST.

Queue A is typically used by arbitrageurs when the price of UST falls below its peg, aiming to generate profit. During the Terra-Luna collapse, the excessive use of this queue led to the exponential minting of LUNA, ultimately causing its price to plummet. The idea is to prioritize the activation of queue B – which triggers the burning of LUNA – by using a probability function, particularly when the price of UST falls below one dollar. This approach could incentivize users to burn LUNA in exchange for UST, as their transaction would likely be processed faster. In contrast, transactions that mint LUNA are placed in queue A, where they may experience longer processing times. Note that when some users choose to burn their own LUNA to mint UST during a depeg event, this action does not result in a profit: for every one dollar's worth of LUNA they burn, they receive one UST – which, at that moment, is worth less than one dollar. However, users may still pursue this strategy in an effort to minimize their losses.

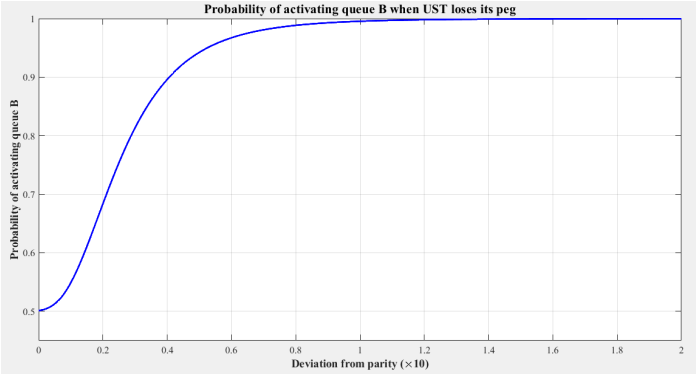

Figure 4: Probability of activating queue B in the algorithmic Queue System stabilization proposal.

In this case as well, the choice of which queue to activate at each iteration when UST loses its peg is governed by a probability function, which determines the probability of activating queue B. It is defined as follows:

$$P_B(x) = e^{-\frac{1}{(x+\beta)^\kappa}} \cdot 0.5 + 0.5 \tag{12}$$

where κ determines the slope of the curve and β is a parameter that adjusts the horizontal position of the function, shifting it left or right. The multiplication by 0.5 and the addition of 0.5 in the function are essential to ensure that the codomain is confined between 0.5 and 1.

The independent variable x represents the deviation from parity, indicating how much the price of UST diverges from one dollar, multiplied by a factor equal to 10. For example, if UST is worth 0.99 dollars, then $x = 0.1$. The probability of activating queue B increases as the peg loss becomes more significant. The shape of this function is illustrated in Figure 4, with $\beta = 0.8$ and $\kappa = 8$.

4.3. Implementing a USDT Reserve Pool

This improvement involves the implementation of a reserve pool that the protocol can use to automatically buy back T_s, i.e., UST, if its price falls below the peg. The reserve pool functions as a collateral for the token T_s. The purpose is to utilize this pool during a crisis to quickly restore $P_s(n)$ to the peg. The reserve pool is filled with R USDT and T_s in equal quantities. In our simulation, we have set

$R = 0.10, 0.20, 0.30 \cdot$ (total supply T_s). If $P_s(n)$ falls below a threshold value Tr_{low}, the protocol uses some of the reserves to buy T_s from Π^S. This mechanism ensures that T_s can efficiently recover the peg value even in adverse market conditions, at least as far as the pool has tokens to use. At iteration n, the quantity x of USDT to sell is determined by the following system of equations, derived from formulas (4) and (5), where $\widetilde{Q}_s(n)$ is the quantity of T_s when $P_s(n) = \mathrm{Tr}_{low}$:

$$\begin{cases} x(n) = \frac{k}{\widetilde{Q}_s(n)} - Q_U(n) \\ \mathrm{Tr}_{low} = \frac{k}{\widetilde{Q}_s(n)^2} \end{cases} \implies \begin{cases} x(n) = \frac{k}{\widetilde{Q}_s(n)} - Q_U(n) \\ \widetilde{Q}_s(n) = \sqrt{\frac{k}{\mathrm{Tr}_{low}}} \end{cases}$$

Moreover, the reserve pool can also intervene when the stablecoin price exceeds the peg, i.e., when $P_s > \mathrm{Tr}_{high}$, by selling a portion of its T_s holdings in exchange for USDT. This sale increases the supply of T_s in Π^S, helping to bring the price back to the peg. The amount of T_s to be sold in this scenario can be calculated by using a similar approach as in the previous case where the price falls below the peg:

$$\begin{cases} x(n) = \widetilde{Q}_s(n) - Q_s(n) \\ \widetilde{Q}_s(n) = \sqrt{\frac{k}{\mathrm{Tr}_{high}}} \end{cases}$$

In this case, $\widetilde{Q}_s(n)$ is the quantity of T_s when $P_s(n) = \mathrm{Tr}_{high}$. The ratio between USDT and T_s in the reserve pool changes dynamically over time as the protocol responds to fluctuations in the stablecoin price. When $P_s < \mathrm{Tr}_{low}$, the reserve pool depletes its USDT to purchase T_s, increasing the amount of T_s held in the pool. Conversely, when $P_s > \mathrm{Tr}_{high}$, the protocol sells T_s to accumulate more USDT, thereby shifting the balance in favor of USDT. This continual adjustment ensures that the protocol has sufficient liquidity to stabilize P_s in both upward and downward market movements. However, this approach also introduces the risk of depleting either USDT or T_s if price deviations persist for an extended period, potentially limiting the effectiveness of the reserve pool in severe market crises. The choice of the two threshold values, Tr_{low} and Tr_{high}, is flexible, provided the following condition holds:

$$Tr_{low} \leq 1 \leq Tr_{high}$$

where 1 represents the target peg value in USD. The specific thresholds determine how sensitive the reserve mechanism is to price fluctuations. A narrower band helps to maintain the peg more precisely but requires a larger reserve buffer to support frequent interventions. In contrast, a wider band conserves reserve resources but may allow the price to drift further from the peg, especially during rapid market shifts. In our simulations, we chose to test two different threshold intervals: the first is $[0.95, 1.00]$ while the second is $[0.98, 1.02]$.

4.4. Implementing a BTC Reserve Pool against LUNA Hyperinflation

This approach is an automatization of the attempt made by the Luna Foundation Guard to support the value of UST, when during the agitated phases of the collapse they sold at the market about 80 000 BTC [29]. The proposed solution aims to prevent the hyperinflation of LUNA, which results from its excessive minting due to the usage pressure of the Terra-Luna protocol of Figure 1. It is well known that the number of LUNA in circulation soared from 340 millions to 6.5 trillions at the end of the collapse event [30].

When the value of the stablecoin T_s reaches a critical level, e.g., \$0.95, a queue system is automatically activated, which involves a reserve of BTC. At this point, a user has three options available:

1. sell their own T_s directly at the market;
2. use the classic stabilization protocol, burning 1 T_s to obtain \$1 worth of the volatile coin T_v;
3. burn 1 T_s to obtain \$1 worth of BTC.

It is necessary to clarify some aspects. For the correct functioning of the system, the presence of a reliable oracle, which provides the real-time price of BTC, is essential. Moreover, users will not receive real BTC, as the latter is a native cryptocurrency of another blockchain. Instead, they will receive \$1 worth of *wrapped* BTC (wBTC), which is a synthetic token that represents the ownership of a BTC on a blockchain different from the original one. Obviously, it is not possible to prevent users from selling their own AS on the market. However, it is possible to induce them to undertake option 2 or option 3 through the creation of a double-queue system: queue A collects transactions of users eager to exchange T_s tokens for BTC, while queue B contains transactions of those who prefer to normally use the protocol, obtaining \$1 worth of T_v for each T_s inserted.

Both options offer the benefit of reducing the circulating supply of T_s, but queue A has the additional advantage of avoiding the minting of T_v. A probability function can govern the mechanism by selecting, at each time interval, which queue to activate. Both queues follow the FIFO (First In First Out) principle. The probability function can be a simple logistic function, which essentially depends on two parameters:

- the deviation of the price $P_s(n)$ from the parity;
- the filling percentage of the BTC reserve.

In our simulation, the function $P_A(x)$ returns the probability that queue A is activated. It is defined as follows:

$$P_A(x) = \frac{1}{1 + e^{-\kappa(x-\alpha)}} \tag{13}$$

where κ determines the slope of the curve and α establishes the point on the x-axis where the graph has an inflection point. These parameters can be manipulated to obtain a function that conforms to one's preference. The independent variable x could be set as a combination of the filling percentage of the BTC reserve and the deviation from parity, e.g., the product of the two. To reproduce the market context before the collapse of the Terra-Luna system, the 100% of the reserve could be considered equal to \$1B in BTC. When the reserve is empty, the probability of queue A being selected is set to zero.

5. Simulation Results

We carried out two distinct types of simulations, each differing in the approach used to model the collapse dynamics.

The first type of simulation applied the variable mean approach based on price of Section 3.5.1. This method focuses on evaluating the robustness of the AS price stabilization improvements. Specifically, it assesses the system ability to preserve the stablecoin peg during volatile conditions. We first examine the original Terra protocol under normal market conditions (Section 5.1), characterized by constant volatility (σ_Δ), while varying its redemption capacities. This allows us to fine-tune simulation parameters for the next phase: we simulate a turbulent market environment with increasing volatility to induce a collapse scenario, aiming to evaluate the effectiveness of the improved VLP replenishing algorithm (Section 5.2) and the automatic UST reserve pool (Section 5.3) in maintaining system stability.

The second simulation type implemented the variable probability approach of Section 3.5.2. It analyzes the effectiveness of the BTC reserve pool (Section 5.4) and a random queue system (Section 5.5) in preventing the hyperinflation of GT. Regarding the BTC reserve, we explore its utility during periods of extreme market stress, providing insights into how such a reserve can contribute to the overall stability of the system. Then we analyze the impact of the purely algorithmic queue system, which prioritizes the minting of UST over LUNA during a crisis scenario.

5.1. Impact of System Parameters on Stablecoin Stability

Firstly, we conducted a series of 100 simulations, each consisting of 50 000 iterations, with varying constant values for the parameters σ_Δ, *PoolRecoveryPeriod* (*RP*), and $Pool_{Base}$. This resulted in a comprehensive total of 2 700 simulation runs. The parameter σ_Δ plays a pivotal role in determining market volatility, while *PoolRecoveryPeriod* and $Pool_{Base}$ determine the redemption capacity of the stability mechanism. In each simulation, T_s starts at \$1, with a total supply $Supply_{T_s}$ of $1.8542 \cdot 10^{10}$. The initial price of the volatile token is set to $P_v(0) = \$78.70$, with an initial supply $Supply_{T_v}$ of $3.44 \cdot 10^8$ units. The selected supplies and prices of T_s and T_v correspond to the actual UST and LUNA supplies as of May 1, 2022 [28]. In contrast, Π^S and Π^V initially consist of $Supply_{T_s} \cdot 0.7\, T_s$ and $Supply_{T_v} \cdot 0.7\, T_v$ respectively.

An event of depegging is set to happen when the price of the stablecoin (P_s) deviates by more than 5% from its peg, meaning that the price falls below 0.95 or goes above 1.05. If the price of the stablecoin falls below 0.90 and, at that point, the stablecoin market cap is higher than the volatile token market cap, we consider the system to have collapsed. Table 3 summarizes the frequency of depegging and collapse events across different parameter configurations. Figure 6 illustrates an example of price behavior of T_s, T_v, and VLP δ variation during a simulation.

To assess the system stability under these scenarios, we computed the *mean squared error* (MSE) relative to a constant unitary function, which reflects the extent of price deviations from the T_s peg. For each parameter combination, we recorded the MSE in all 100 runs, then report in Figure 5 both the average MSE and its 95% confidence interval. This approach allows us to evaluate the typical price deviation and its variability across repeated simulations.

In the low-volatility case ($\sigma_\Delta = 10^{-5}$, Figure 5a), the average MSE remains extremely low (on the order of 10^{-7}) for all recovery periods and pool sizes, indicating that the original protocol can easily face such volatility and almost never depegs. However, for the slowest replenishment setting ($RP = 5\,000$), the confidence interval widens noticeably: a small subset of runs experiences transient depegs that drive up the squared error, even though the majority of runs maintain peg stability. This variability highlights that, under very low volatility, large RP can still permit occasional outliers.

At moderate volatility ($\sigma_\Delta = 10^{-4}$, Figure 5b), average MSE values increase slightly (up to 10^{-3}), and the influence of pool size becomes clearer. Shorter recovery periods ($RP = 50,\ 500$) produce consistently tight confidence bands around near-zero MSE, whereas $RP = 5\,000$ again shows a broad interval, reflecting that slower replenishment sometimes fails to correct price deviations before they amplify, even if most runs succeed in re-pegging.

Under high volatility ($\sigma_\Delta = 10^{-3}$, Figure 5c), all configurations exhibit substantial MSE (order unity and above), and confidence intervals grow large, signifying high run-to-run variability in collapse timing and severity. Yet, the trend is clear: increasing $Pool_{Base}$ and reducing *RP* both lower the average MSE and tighten its confidence interval, confirming that greater liquidity and faster recovery robustly mitigate peg deviations in turbulent markets.

5.2. Effectiveness of the Improved TerraPool_δ *Replenishing Mechanism in Maintaining Price Stability*

In the second series of tests, we conducted 100 simulations with the goal of inducing the system to collapse by gradually increasing the volatility σ_Δ from $5.00 \cdot 10^{-9}$ to $1.00 \cdot 10^{-3}$, with a $5.00 \cdot 10^{-9}$ step every iteration and a total of 200 000 iterations. The results are shown in Table 4 and make it evident that the original implementation of Terra is unable to withstand such scenarios, with the peg being lost in every simulation and a collapse occurring in 71 out of 100 simulations. On the other hand, the modified $TerraPool_\delta$ replenishing mechanism yielded better outcomes, with only 28 collapses and a significant improvement in the overall stability of the system. This is evident from the reduction in depeg events (from 100 out of 100 in the original protocol to 78 out of 100 with the proposed improvement) and the decrease in the average MSE (from 0.070 to 0.055).

A two-sample t-test comparing the MSE distributions of the improved mechanism against the original yields a p-value of 0.267 (Welch's t-test), indicating that this reduction in MSE is not statistically significant at the 0.05 level.

5.3. Effectiveness of an Automatic USDT Reserve Pool in Maintaining Price Stability

The partial collateralization of UST, with a threshold interval of [0.95, 1.00], resulted in only a slight reduction in the number

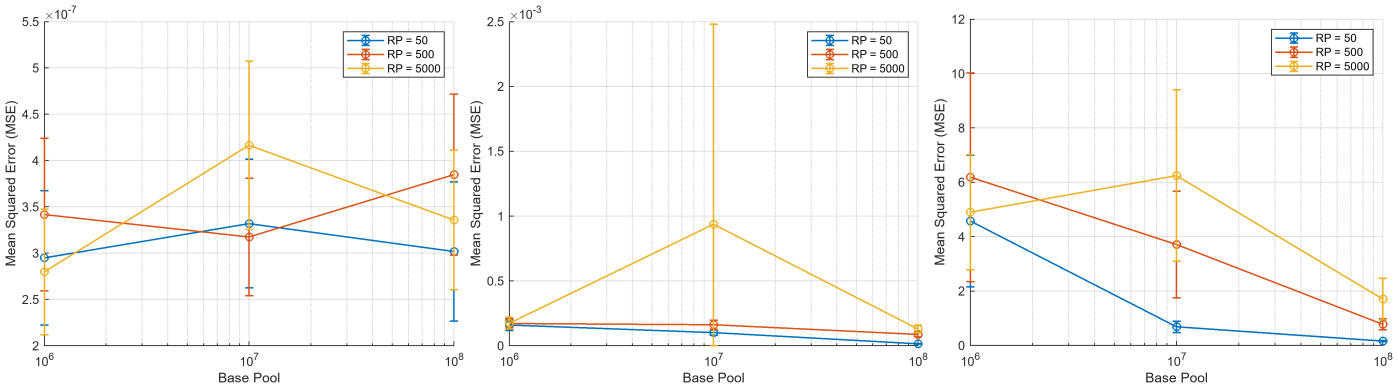

(a) $\sigma = 10^{-5}$ (b) $\sigma = 10^{-4}$ (c) $\sigma = 10^{-3}$

Figure 5: Mean squared error (MSE) across varying base pool sizes and recovery periods, with 95% confidence intervals over 100 runs. Panels (a)–(c) correspond to volatility levels $\sigma_\Delta = 10^{-5}, 10^{-4}, 10^{-3}$, respectively. At low volatility ((a),(b)), average MSE remains near zero but exhibits high variance for slow replenishment (RP=5000), indicating occasional depeg outliers. Under high volatility (c), larger base pools and shorter recovery periods consistently lower both MSE and its spread, demonstrating clearer stabilizing effects.

Table 3: Depeg and collapse rates of the algorithmic stablecoin T_s under varying market volatility (σ_Δ), VLP recovery periods (RP), and base pool sizes. Each cell reports the percentage of simulations (out of 100) in which T_s lost its peg ("depeg") and in which the system collapsed, along with 95% Clopper–Pearson confidence intervals. Results show that under very low volatility ($\sigma = 10^{-5}$), no depegs or collapses occur, whereas high volatility ($\sigma = 10^{-3}$) yields the collapse rates between 54% and 72% depending on RP and pool size. Intermediate volatility ($\sigma = 10^{-4}$) produces sporadic depeg and collapse events only for long recovery periods (RP=5000) with low base pools.

RP	$Pool_{Base}$	$\sigma = 10^{-5}$	$\sigma = 10^{-4}$	$\sigma = 10^{-3}$
50	10^6	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	2.00% [0.24,7.04] | 0.00% [0.00,3.62]	100.00% [96.38,100.00] | 66.00% [55.85,75.18]
	10^7	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	100.00% [96.38,100.00] | 59.00% [48.71,68.74]
	10^8	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	100.00% [96.38,100.00] | 69.00% [58.97,77.87]
500	10^6	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	5.00% [1.64,11.28] | 0.00% [0.00,3.62]	100.00% [96.38,100.00] | 59.00% [48.71,68.74]
	10^7	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	3.00% [0.62,8.52] | 0.00% [0.00,3.62]	100.00% [96.38,100.00] | 64.00% [53.79,73.36]
	10^8	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	100.00% [96.38,100.00] | 61.00% [50.73,70.60]
5000	10^6	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	6.00% [2.23,12.60] | 0.00% [0.00,3.62]	100.00% [96.38,100.00] | 57.00% [46.71,66.86]
	10^7	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	5.00% [1.64,11.28] | 1.00% [0.03,5.45]	100.00% [96.38,100.00] | 54.00% [43.74,64.02]
	10^8	0.00% [0.00,3.62] | 0.00% [0.00,3.62]	1.00% [0.03,5.45] | 0.00% [0.00,3.62]	100.00% [96.38,100.00] | 72.00% [62.13,80.52]

of collapses compared to the original implementation. Furthermore, as shown in Table 4, there is no clear increase in the stability of the system as we enhance the collateral coverage. This unexpected behavior is related to the threshold interval choice. Since the UST reserve pool automatically sells UST when P_s goes above \$1, the stabilization mechanism is prevented from minting new stablecoins and, consequently, burning volatile tokens. This implies that the volatile token supply tends to increase over time, reducing the token price P_v and compromising the effectiveness of the stabilization mechanism. As a consequence, we chose to try another threshold interval for the UST reserve pool intervention: [0.98, 1.02]. This symmetric interval reduced the likelihood of early UST sales at prices just above 1, which mitigates the previous issue where the minting of new stablecoins was restricted. This, in turn, prevents unchecked growth of the volatile token supply, helping to stabilize both the volatile token price P_v and the overall system.

Finally, we combined the improved $TerraPool_\delta$ replenishing mechanism with the partial collateralization approach by using the adjusted intervention threshold of [0.98, 1.02]. This combination aimed to leverage the benefits of both mechanisms: maintaining stability through targeted reserve interventions while avoiding excessive volatile token depreciation. By utilizing our improved VLP replenishing algorithm and implementing a partial collateralization of only 30% of the total T_s supply, we achieved a significant reduction in the number of collapses (−95.8%) and in the average MSE across the simulation set (−94.9%) compared to the original protocol. These results are summarized in the final row of Table 4.

We performed Welch's two-sample t-tests on the MSE distributions of each scenario against the original protocol. For the reserve-pool intervention alone with thresholds $\text{Tr}_{low} = 0.98, \text{Tr}_{high} = 1.02$, the p-values were 0.714, 0.664, and 0.00443 for 20%, 40%, and 60% collateralization, respectively. When combining the VLP improvement with partial collateralization, the p-values fell to 0.0101 (20%), 0.000195 (40%), and 0.000179 (60%). At the $\alpha = 0.05$ level, only the 60% reserve alone and all combined interventions achieve statistical significance, confirming that higher reserve coverage and hybrid designs yield statistically significant improvements in price stability over the original system.

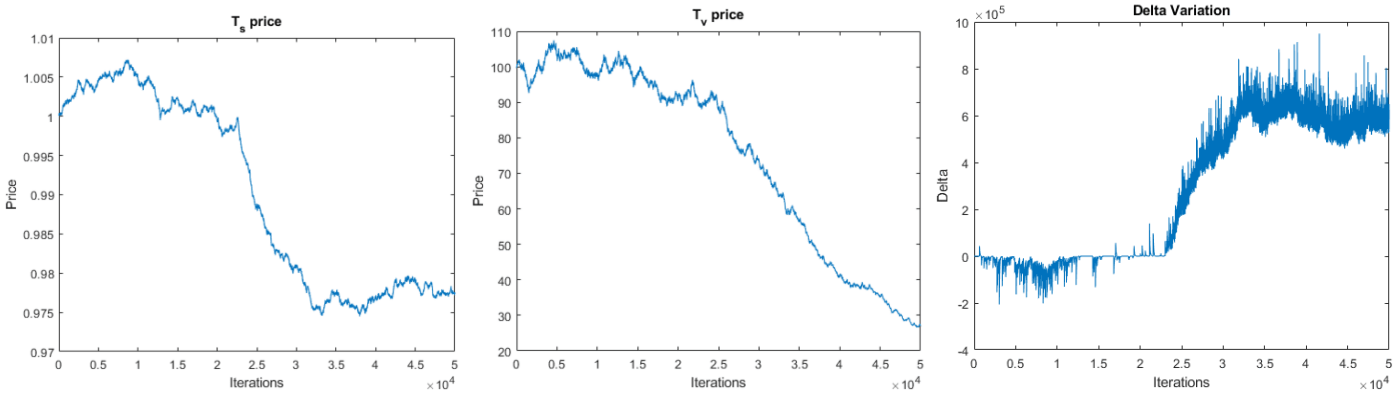

(a) T_s price across iterations. (b) T_v price across iterations. (c) δ value across iterations.

Figure 6: Comparison of the prices of T_s, T_v, and the VLP δ variations across iterations.

5.4. *Effectiveness of a BTC Reserve Pool in Preventing Hyperinflation*

As for the BTC Reserve Pool improvement, we performed two separate sets of tests. The first one involved the use of the Terra standard stabilization protocol, without any queue system. The second one incorporated the queue system described in Section 4.4. In both scenarios, we conducted 100 runs of the simulation program, monitoring and recording the circulating supply of T_v (LUNA). The results are shown in Figure 7.

5.5. *Effectiveness of the Stochastic Queue System in Preventing LUNA Hyperinflation*

To test the effectiveness of the stochastic queue system we perform two kinds of simulations: the first one with the original Terra-Luna protocol and the second one with the implementation of the stochastic queue system described in Section 4.2.

The introduction of a probability function that promotes the burning of LUNA during a peg loss effectively achieves its goal. In fact, while the supply of LUNA dramatically explodes to the order of 10^{12} after only 400 time instants using the standard protocol (see Figure 8), the supply of LUNA even decreases when we introduce the stochastic queue system, as illustrated in Figure 9. Avoiding the hyperinflation of LUNA helps to prevent its value from collapsing to zero; however, its value is still decreasing, albeit at a slower rate. The price trend of LUNA with the random queue system is illustrated in Figure 10. The other side of the coin is that prioritizing queue B – which collects transactions from those who want to burn their LUNA to obtain UST – leads to an increased minting of UST, which augments its circulating supply and consequently decreases its value even more rapidly. Moreover, once arbitrageurs burn LUNA to obtain UST, they are likely to sell the newly obtained UST, leading to a further decline in the price of the stablecoin.

However, the circulating supply of UST (see Figure 11) does not increase as dramatically as that of LUNA, which reached levels in the trillions. Such a high supply could represent a point of no return and the fact that this solution does not reach such quantities may be an advantage. In our simulation the circulating supply of UST increases by an average of only 0.15% when implementing the queue system, a negligible amount compared

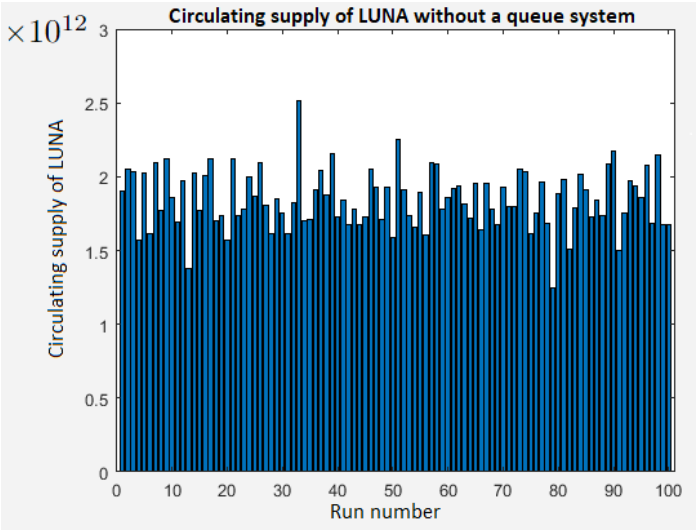

(a) Circulating supply of LUNA in the original protocol.

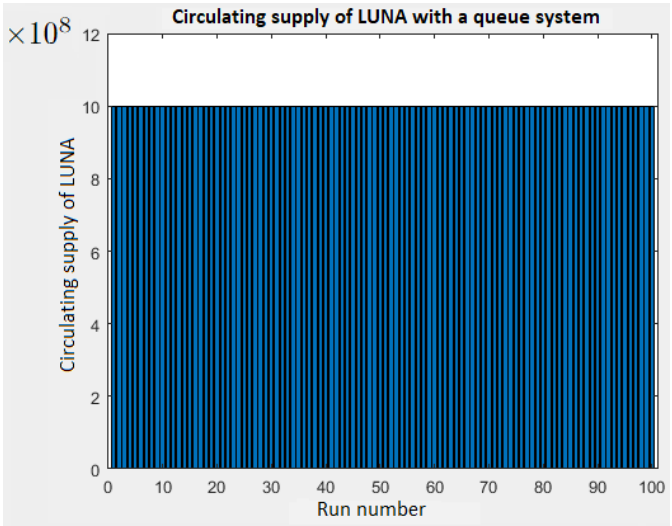

(b) Circulating supply of LUNA with the queue system implementing the BTC Reserve Pool.

Figure 7: Supply of LUNA without and with the queue system described in Section 4.4.

Simulation	Total collapses	Depeg events	Mean collapse time	MSE
Original	71	100	126771	0.070
Improved *TerraPool*$_\delta$ replenishing	28	78	145437	0.055
UST Reserve Pool 20% (**Tr**$_{low}$ = 0.95, **Tr**$_{high}$ = 1.00)	58	100	134641	0.066
UST Reserve Pool 40% (**Tr**$_{low}$ = 0.95, **Tr**$_{high}$ = 1.00)	57	100	143200	0.065
UST Reserve Pool 60% (**Tr**$_{low}$ = 0.95, **Tr**$_{high}$ = 1.00)	62	100	141751	0.040
UST Reserve Pool 20% (**Tr**$_{low}$ = 0.98, **Tr**$_{high}$ = 1.02)	40	98	153008	0.033
UST Reserve Pool 40% (**Tr**$_{low}$ = 0.98, **Tr**$_{high}$ = 1.02)	29	82	163692	0.025
UST Reserve Pool 60% (**Tr**$_{low}$ = 0.98, **Tr**$_{high}$ = 1.02)	24	59	175342	0.016
Improved *TerraPool*$_\delta$ replenishing + UST Reserve Pool 20% (**Tr**$_{low}$ = 0.98, **Tr**$_{high}$ = 1.02)	13	63	157313	0.020
Improved *TerraPool*$_\delta$ replenishing + UST Reserve Pool 40% (**Tr**$_{low}$ = 0.98, **Tr**$_{high}$ = 1.02)	19	50	175399	0.013
Improved *TerraPool*$_\delta$ replenishing + UST Reserve Pool 60% (**Tr**$_{low}$ = 0.98, **Tr**$_{high}$ = 1.02)	3	32	178411	0.0036

Table 4: Number of collapses across 100 simulations, with increasing value of the volatility σ_Δ, for the original protocol and the first two improvement proposals. Mean collapse time refers to the average iteration when the system collapsed (i.e., when $P_s < 0.9$ and T_v capitalization $< T_s$ capitalization). Price stability is measured using the MSE. The terms Tr_{low} and Tr_{high} represent the lower and upper bounds of the reserve interval intervention. Market dynamics are modeled by using the variable mean approach based on price (Section 3.5.1).

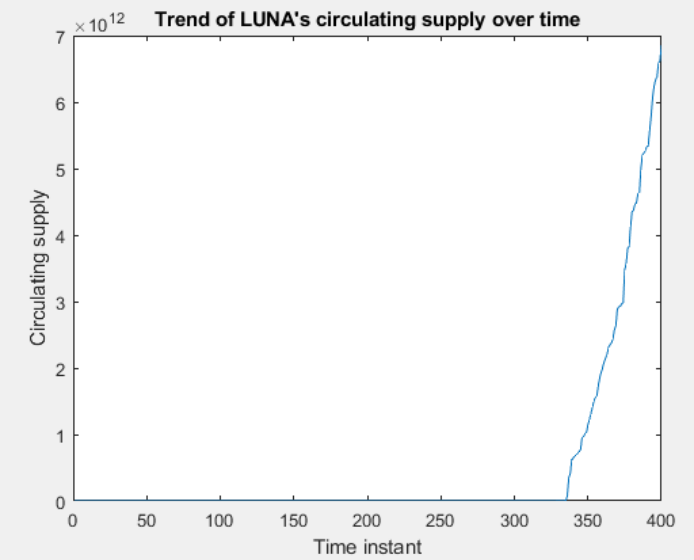

Figure 8: Circulating supply of LUNA during the collapse.

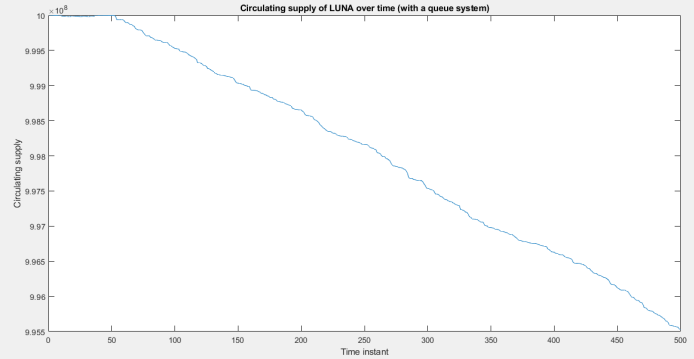

Figure 9: Circulating supply of LUNA during the collapse with queue system.

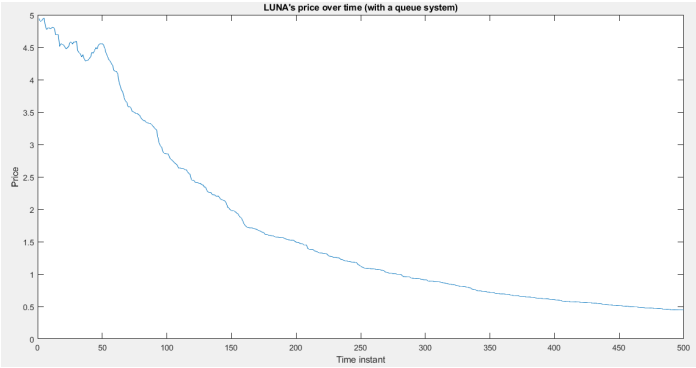

Figure 10: Price of LUNA during the collapse with queue system.

to the 1 900 000% increase in the supply of LUNA during the real collapse.

At this point, one might wonder why the increase in UST supply in the simulation is so restricted, despite the queue system incentivizing UST minting during a crisis scenario. The reason lies in the mechanism that allows one UST to be minted by inputting one dollar's worth of LUNA into the system. During a collapse, a dollar's worth of LUNA might represent a large number of LUNA tokens if its price is plummeting, thus requiring the burning of many LUNA to mint just one UST.
In reality, however, the opposite occurred more frequently: the queue system was predominantly used to burn UST in exchange for minting one dollar's worth of LUNA for each burnt UST. In this scenario, with LUNA's price in free fall, burning even a single UST resulted in minting an enormous quantity of LUNA to reach the equivalent value of one dollar.

Nevertheless, it is challenging to raise the price of UST without demand from users. Not reaching such supply levels may still provide ecosystem designers with the time needed to implement other solutions, such as scheduled UST purchases.

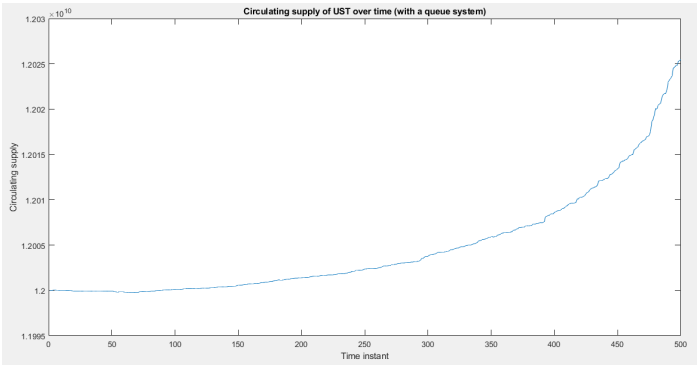

Figure 11: Supply of UST during the collapse with queue system.

6. Discussion

The first conclusion we can draw is the structural weakness of the Terra-Luna protocol, since Table 4 shows an astonishingly high number of collapses as a consequence of an increasing volatility σ_Δ of the market (71 out of 100). Also, the huge number of circulating supply of LUNA, of the order of 10^{12} as illustrated in Figure 7a, is a clear warning of this structural weakness. Even if a perfect AS should resist all kinds of market destabilization forces, with only "0" in the second and third columns of Table 3, the improvements we obtain with the four proposals shed light on the mechanism one could implement to increase the strength of an AS protocol.

The *TerraPool*$_\delta$ improvement indicates that it is possible to work on the replenishment protocol with a good rate of success, which results in a reduction of 60.6% in total collapses (28 collapses out of 100 instead of 71). Using an automatic USDT reserve pool can help to stabilize the system, but it is important to choose correctly the activation threshold interval $[\mathrm{Tr}_{low}, \mathrm{Tr}_{high}]$. Implementing a symmetric interval can lead to good results, especially for larger reserves where we observe a 66.2% reduction in the number of collapses. (40, 29, and 24 collapses out of 100 collateralizing T_s to respectively 10%, 20%, and 30% of its total supply). Combining these two techniques can result in a highly resilient system capable of withstanding extreme market volatility. Specifically, our analysis of the improved *TerraPool*$_\delta$ replenishment mechanism, along with a partial collateralization strategy covering 30% of the total stablecoin supply, reveals that only 3 out of 100 collapses occurred, resulting in a reduction of 95.8% in total collapses. These failures were not directly caused by the stablecoin's mechanics but resulted from a steep decline in the market capitalization of T_v due to extreme market conditions. This capitalization drop preceded the stablecoin's collapse, highlighting that external market forces, rather than flaws in the stablecoin protocol itself, were responsible.

Also, the method of limiting the LUNA supply with the queue system implementing the BTC Reserve Pool shows a good performance, reducing to 10^9 the original supply of LUNA of order 10^{12}. Note that this huge supply was the main cause of the crash in the LUNA value and of the Terra-Luna collapse. Note, moreover, that this last method guarantees an almost constant value of supply for all simulations. These evaluations underscore the critical importance of incorporating reserve pools and refining virtual pool replenishing mechanisms to improve the stability of AS systems in the face of market volatility. The utilization of a reserve pool provides a significant cushion against adverse market movements, acting as a stabilizing force during periods of heightened volatility. The obtained data serve as a clear illustration of the robustness exhibited by this hybrid collateralized-AS. It effectively demonstrates the coin's ability to resist challenges posed by a highly volatile market. In general, our solutions extend beyond Terra-Luna, adapting to diverse seigniorage stablecoin frameworks.

As far as the limitations of our approach are concerned, while the introduction of a reserve pool shows promising results, it moves away from the notion of a pure AS, transitioning towards a partially collateralized stablecoin model. An attempt to enhance the resilience of the Terra-Luna ecosystem while remaining within a purely algorithmic framework is found in the fourth proposal. The introduced queue system, which prioritizes the burning of LUNA during de-pegging, effectively helps to prevent hyperinflation of both LUNA and UST, as their circulating supply does not increase significantly. However, the value of UST tends to decline more rapidly due to the increase in its circulating supply and the subsequent sales. Nevertheless, the supply of both tokens remains limited, which could be an advantage compared to the original approach.

Future research should address these trade-offs and explore avenues for optimizing stability while maintaining algorithmic integrity. Another limitation is related to the design choices we made. They are an inevitable approximation of the human behavior in financial markets. Anyway, we tried to be as general as possible in designing these phenomena, since our simulations are initializable with different parameters, allowing for flexibility in capturing various market conditions and behaviors. A critical aspect for future explorations lies in refining the parameters and assumptions underlying our simulation model. This could be achieved by incorporating real-world data and historical market trends.

Moving forward, our research suggests several interesting directions. Further investigation into alternative replenishment protocols and reserve pool mechanisms could yield insights into improving the stability of this type of system. Our simulation-based approach could help researchers to effectively design and evaluate new algorithmic stablecoin protocols, offering valuable information about their performance and robustness in various market conditions.

7. Conclusions

This study has investigated the collapse dynamics of the Terra-Luna ecosystem, emphasizing the role and the structural weakness of its stabilization mechanisms. We have employed a simulation environment to model the Terra-Luna collapse that occurred in May 2022, replicating both market transactions and protocol-specific activities. By carefully analyzing the involved parameters, we have been able to recreate the conditions that led to the deviation from the UST peg and its subsequent collapse. From our simulations, it is evident that the original Terra-

Luna stabilization protocol suffered from intrinsic weaknesses, as shown by an alarmingly high rate of collapse under increasing market volatility. However, by modifying the original Terra protocol, we have demonstrated potential avenues to significantly improve the stability of algorithmic stablecoins.

Acknowledgments. This research has been supported by the PRIN 2020 project *NiRvAna – Noninterference and Reversibility Analysis in Private Blockchains*. We extend our gratitude to BAX for providing access to its computational resources, which were crucial for running the simulation environment. We acknowledge the financial support from the PNRR (Piano Nazionale di Ripresa e Resilienza) fund.

Code Repository. The Matlab code that implements our simulations is publicly available at https://github.com/FedericoCalandra/Algorithmic-Stablecoin-Collapse-Simulation.

References

[1] K. Qin, L. Zhou, Y. Afonin, L. Lazzaretti, A. Gervais, CeFi vs. DeFi — Comparing Centralized to Decentralized Finance, arXiv preprint arXiv:2106.08157 (2021).

[2] Ampleforth, #ampl the decentralized unit of account, https://www.ampleforth.org/, accessed: 2024-10-17 (2022).

[3] E. Kereiakes, D. Kwon, M. Di Maggio, N. Platias, Terra money: Stability and adoption, White Paper (2019).

[4] Anchor protocol, accessed: 2024-10-19.
URL https://docs.anchorprotocol.com/anchor-2

[5] H. Uhlig, A Luna-tic stablecoin crash, Tech. rep., National Bureau of Economic Research (2022).

[6] F. Calandra, F. P. Rossi, F. Fabris, M. Bernardo, Making algorithmic stablecoins more stable: The Terra-Luna case study, in: Proceedings of the 6th Distributed Ledger Technology Workshop (DLT 2024), Vol. 3791 of CEUR Workshop Proceedings, CEUR-WS.org, 2024, pp. 19:1–19:14.

[7] V. Mohan, Automated market makers and decentralized exchanges: A DeFi primer, Financial Innovation 8 (1) (2022) 20.

[8] Y. Zhang, X. Chen, D. Park, Formal specification of constant product (xy= k) market maker model and implementation, White paper (2018).

[9] terra.money, Liquidity parameters, https://classic-agora.terra.money/t/liquidity-parameters-3/3895, accessed: 2024-10-17 (2022).

[10] A. Briola, D. Vidal-Tomás, Y. Wang, T. Aste, Anatomy of a stablecoin's failure: The Terra-Luna case, Finance Research Letters 51 (2023) 103358.

[11] J. Liu, I. Makarov, A. Schoar, Anatomy of a run: The Terra luna crash, Tech. rep., National Bureau of Economic Research (2023).

[12] F. van Echelpoel, M. T. Chimienti, M. Adachi, P. Athanassiou, I. Balteanu, T. Barkias, I. Ganoulis, D. Kedan, H. Neuhaus, A. Pawlikowski, et al., Stablecoins: Implications for monetary policy, financial stability, market infrastructure and payments, and banking supervision in the Euro area, ECB Occasional Paper, European Central Bank 247 (2020).

[13] C. Calcaterra, W. A. Kaal, V. Rao, Stable cryptocurrencies: First order principles, Stanford Journal of Blockchain Law & Policy 3 (2020) 62–64.

[14] A. Kosse, M. Glowka, I. Mattei, T. Rice, Will the real stablecoin please stand up?, BIS Papers, Bank for International Settlements 141 (2023).

[15] L. Ante, I. Fiedler, J. M. Willruth, F. Steinmetz, A systematic literature review of empirical research on stablecoins, FinTech 2 (1) (2023) 34–47.

[16] J. Choi, H. H. Kim, Stablecoins and central bank digital currency: Challenges and opportunities, Oxford Research Encyclopedia of Economics and Finance (March 2024). doi:10.2139/ssrn.4756822.

[17] R. Clements, Built to fail: The inherent fragility of algorithmic stablecoins, SSRN (October 2021).

[18] J. Cho, A token economics explanation for the de-pegging of the algorithmic stablecoin: Analysis of the case of Terra, Ledger 8 (2023) 36–44.

[19] G. Kurovskiy, N. Rostova, How algorithmic stablecoins fail, Tech. rep., SNB BNS (2023).

[20] S. Ferretti, M. Furini, Cryptocurrency turmoil: Unraveling the collapse of a unified stablecoin (USTC) through Twitter as a passive sensor, Sensors 24 (2024) 1270. doi:10.3390/s24041270.

[21] A. Adams, M. Ibert, Runs on algorithmic stablecoins: Evidence from iron, titan, and steel, FEDS Notes (2022).
URL https://www.federalreserve.gov/econres/notes/feds-notes/runs-on-algorithmic-stablecoins-evidence-from-iron-titan-and-steel-20220602.html

[22] K. Saengchote, K. Samphantharak, Digital money creation and algorithmic stablecoin run, SSRNAvailable at SSRN: http://dx.doi.org/10.2139/ssrn.4710515 (2024).

[23] S. Krogstrup, T. Sangill, M. von Sicard, Containing technology-driven bank runs, IMF Blog (2024).

[24] Uniswap, Pricing in uniswap, https://docs.uniswap.org/sdk/v2/guides/pricing, accessed: 2024-10-17 (2018).

[25] Blockchair Wallet Data, Blockchair, https://blockchair.com/, accessed: 2024-10-17 (2023).

[26] BitInfoChart, Bitcoin rich list, https://bitinfocharts.com/top-100-richest-bitcoin-addresses.html, accessed: 2024-10-17 (2024).

[27] G. F. Lawler, V. Limic, Random Walk: A Modern Introduction, Vol. 123, Cambridge University Press, 2010.

[28] TradingView, Cryptocurrency market data: Ust and luna prices and supply on may 1, 2022, https://www.tradingview.com, accessed: 2024-10-17.

[29] Coindesk, Lfg reserves dwindle to just 313 bitcoins from 80k after UST crash, https://www.coindesk.com/business/2022/05/16/luna-foundation-guard-left-with-313-bitcoin-after-ust-crash, accessed: 2024-10-17 (2022).

[30] T. Block, Luna supply soared to 6.5 trillion coins before Terra's latest halt, https://www.theblock.co/post/146762/luna-supply-soared-to-6-5-trillion-coins-before-terras-latest-halt.html, accessed: 2024-10-17 (2022).