On Aave's E-mode Classes and their Risk Parameters

Chaos Labs

March 2023

Abstract

Aave V3 introduces Efficiency mode (E-mode), whereby the protocol can enable more capital-efficient borrowing for highly correlated assets. E-mode is currently being considered for stablecoins pegged to the U.S. dollar, and ETH liquid staking derivatives. We argue that E-mode asset classes must exhibit pairwise mean-reversion properties with suitable mean reversion speeds to prevent the accrual of bad debt to the protocol. This mean-reversion framework can be used, for example, to complement Agent Based Simulation engines for more accurate pricing trajectories. We design a preliminary framework to set E-mode liquidation thresholds (LT), such that they are unlikely to be abused by adversarial agents during moments of price volatility. We also propose loan to value parameters such that retail traders - those not actively managing their health factors - are less exposed to potential liquidations. Finally, we consider a utility function that the protocol can use to choose which assets to include in a particular E-mode class.

1 Introduction

At a high level, the Aave V3 technical specification states that E-mode was "designed to maximize capital efficiency when collateral and loaned assets are correlated in price, particularly when both are derivatives of the same underlying asset" [2]. We find that correlation is not an appropriate test to ensure E-mode assets will achieve higher capital efficiency without increasing value at risk (VaR). Counter-intuitively, assets within similar asset classes (such as stablecoins) can still exhibit strongly negative pairwise price correlations [9]. For instance, USDC and USDT returns are negatively correlated, likely due to the fact that when either token devalues, holders flee to the other token for safety. Conversely, USDC and DAI are positively correlated, likely due to the fact that DAI reserves are materially composed of USDC holdings. We saw both these dynamics play out during the bank run of Silicon Valley Bank in March 2023.

The goal of E-mode is to allow higher liquidation thresholds without incurring materially more bad debt from missed or adverse liquidations. Bad debt is whatever outstanding debt cannot be profitably liquidated using deposited collateral over some period of time. For simplicity, we can consider some outstanding debt as bad debt if:

1. **Missed Liquidation:** Debt can't can't be profitably liquidated within 24 hours (the exact interval is not central to our argument).
2. **Adverse Liquidation:** All the collateral has been liquidated, but some debt remains.

.

Intuitively, the former occurs when the price of the loaned asset rises relative to the price of the collateral asset, *and doesn't converge back to equilibrium within the specified time frame.* The latter occurs if the liquidation threshold is too low, relative to the current (deflated) prices of the assets.

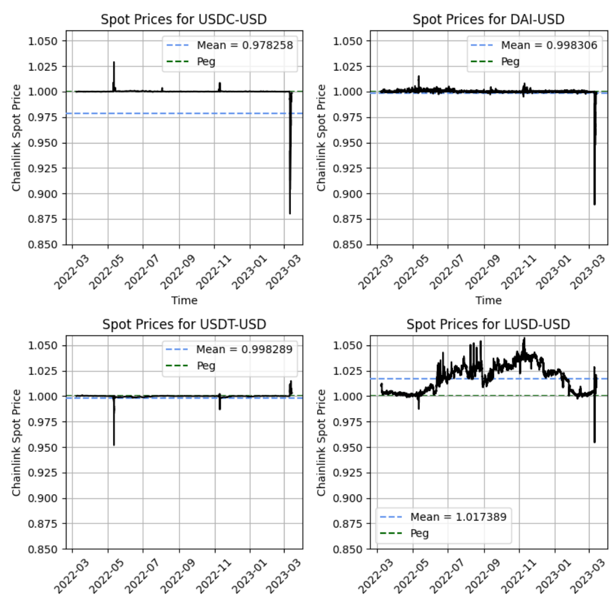

Figure 1: Stablecoin prices from March 2022 to March 14th, 2023.

While positions are underwater - having a health factor below 1 - liquidators can liquidate the deposited lien in order to repay the outstanding loan. Depending on the liquidation bonus and the on-chain liquidity for the asset in question, these liquidations can be unprofitable for the liquidator, in which case there might be a mismatch in assets (liens) and liabilities (loans). This can lead to a gap in Aave's balance sheet, where there isn't enough collateral to account for outstanding loans. Notice, however, that if loaned asset prices fall relative to collateral assets, this gap is closed, and while loaned asset prices fall, liquidations might become profitable.

We argue that E-mode should capture the set of assets for which (a) there is enough on-chain liquidity in all assets such that liquidations incur minimal slippage, and (b) we can be fairly confident that the exchange rate between assets converge to a constant long-term mean. Since (a) can be observed using DEX aggregators such as 1inch, and is/will be accounted for in the risk engines of Aave's risk providers, we will focus on (b) throughout this paper. We find that this property, known in financial engineering as mean-reversion, is very attractive for high efficiency mode assets. If we are fairly confident that the exchange rate of any two assets in an E-mode class will quickly revert to some constant long-term mean, then we can also be confident that we are minimizing potential bad debt.

In this paper, we will discuss the mean reverting properties of derivatives that are pegged to some underlying asset. These include stablecoins pegged to the U.S. dollar and liquid staking derivatives (LSDs), representing future $1:1$ claims on ETH. These assets often exhibit mean-reversion due to guaranteed $1:1$ redemptions by providers such as Circle. Mean-reversion due to high market betas, such as with ETH and BTC, is left for future work. We will test, to some confidence, whether these assets have consistently mean-reverted to an exchange rate of ≈ 1, e.g. USDC will be traded for 1 DAI, or 1 stETH will be exchanged for 1 cbETH. We will measure the time to mean reversion by taking an "epsilon-ball" approach: we create a band around the mean, $[\mu - 0.001, \mu + 0.001]$, and measure the time taken between exiting the band, and re-entering the band. This band is our epsilon-ball, or in this case an epsilon-band, where epsilon is our tolerance $\epsilon = 0.1\%$.

We discuss some important preliminary considerations for setting E-mode risk parameters in Section 2 and our testing framework for detecting mean-reversion and choose E-mode assets in Section 3. We ultimately find that stablecoins such as USDC, USDT, and DAI successfully pass our litmus tests for mean-reversion, whereas liquid staking derivatives don't, likely due to insufficient historical data.

1.1 Counterparty Risks as of March 2023

This paper was written shortly before the Silicon Valley Bank collapse and subsequent momentary de-peg of USDC. This unfortunate event provides insight into why stablecoins exhibit mean-reverting properties relative to the U.S. dollar, and why this mean-reversion is different than that of high market beta mean-reversion, such as the relationship between ETH and BTC. The mean-reversion underpinning stablecoins is driven by the perceived legitimacy of the issuer of said stablecoin, this is the case with Circle's USDC, Tether's USDT, etc.. They might lose their peg if depositors lose faith in the issuer's commitment or ability to redeem tokens at a $1:1$ ratio within a reasonable timeframe. Notice this also applies to the ability of algorithms underpinning algorithmic stablecoins of doing the same.

This claim is consistently tested by the market. Most stablecoins saw sharp outflows as various crypto exchanges and lenders went bankrupt throughout 2022. Tether, for example, saw over $700M$ USDT being redeemed within a 24 hour period following the collapse of FTX in November of 2022 [8]. Despite the spike in redemptions, Tether quickly regained its dollar-peg. Other stablecoins, particularly Terra's algorithmic stablecoin UST, were not as successful.

In this paper, we attempt to formalize, and test to a 99% confidence interval, that derivative assets such as stablecoins or ETH LSDs have successfully maintained their pegs throughout 2022 and early 2023, a particularly volatile period. Further, we test that they regain their pegs at sufficient speeds. However, as discussed in Section 4, we make no attempt at quantifying the risk that historically mean-reverting assets will lose their pegs due to exogenous events. We cannot glean this information from statistical data. As we have seen with the latest USDC de-peg, these events may happen for a variety of unforeseen reasons.

To prevent Aave from being exposed to the fat-tail risk of these assets permanently de-pegging, we consider two possible constraints on E-mode classes:

1. E-mode status should be reserved to assets whose mean-reversion is guaranteed by some counterparty (centralized, or code). Examples of this include Circle's USDC or Lido's stETH.

A more sophisticated risk framework may be developed for other related assets, such as ETH and BTC, but is not addressed by this paper.

2. E-mode status should be reserved to assets whose issuing counterparties are either:

   (a) Regulated and/or audited by relevant financial auditors. This is especially relevant as regulatory scrutiny with respect to stablecoins gains traction in U.S and E.U. courts. See the proposed U.S.'s Stablecoin TRUST act for an example of what this regulation might look like [11]. A few key points to consider are:

   - Disclosures of reserve assets and outstanding liabilities. Refer to Figure 7 in the appendix for a December 2023 audit of Tether's reserves by Binder Dijker Otte (BDO), a globally recognized accounting firm.
   - Regulation on counterparty, duration, liquidity, and market risks of the issuing corporation. For example, Coinbase and Circle are regulated under the New York State Department of Financial Services (NYDFS), as well as several other state regulators [6]

   (b) Decentralized entities with open-sourced code and smart contracts. In this case, the reserves backing up the derivative assets are publicly known and auditable. Such is the case, for example, with Lido's stETH.

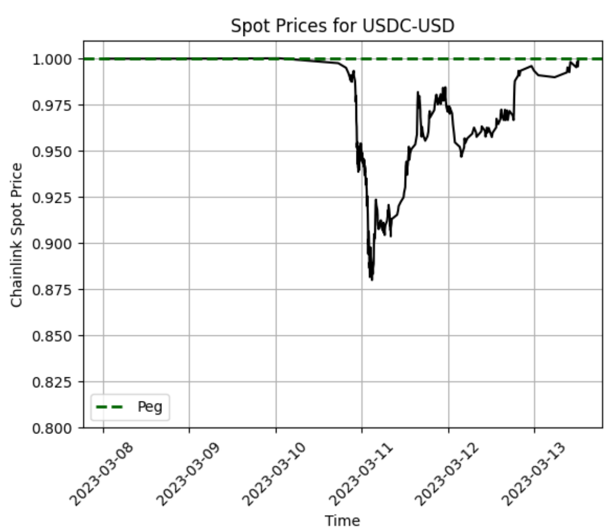

Figure 2: Chainlink spot prices for USDC throughout the Silicon Valley Bank collapse.

Aave's counterparty risk with respect to these issuers has always existed, with or without E-mode. Financial audits and regulation are a step towards mitigating this risk, and we see more issuers complying with such requirements as of 2023. A full review of stablecoin/LSD regulation and audits is outside the scope of this paper. In this paper we assume that the assets in question satisfy these conditions, an assumption we believe was tested for Circle's USDC upon the collapse of Silicon Valley Bank, see Figure 2 for USDC's performance.

2 E-mode Risk Parameters

The liquidation threshold balances capital efficiency and value-at-risk for Aave. We show how the liquidation threshold can be gamed by an adversarial agent, presumably a statistical arbitrage trader, and how a similar concept can be applied to LTV to help retail traders avoid liquidations in high efficiency mode. Here we set LTV and LT as if an E-mode class includes only 2 assets. If the E-mode class contains more than 2 assets, we would consider the lowest LT and LTV observed for any pair in the asset class.

Furthermore, in setting a liquidation threshold according to the highest (observed) deviation in prices, we are implicitly minimizing the risk of adverse liquidations. This will best be measured via agent based simulations, which account for the preferences of liquidators, but we provide an initial conservative estimate.

Finally, we argue the liquidation bonus is primarily a function of on-chain liquidity and we are aligned with the community's current consensus of keeping it in the $1\% - 2\%$ range for stablecoins. We discuss improvements to this methodology, and how it fits into the context of more sophisticated risk engines, in Section 4.

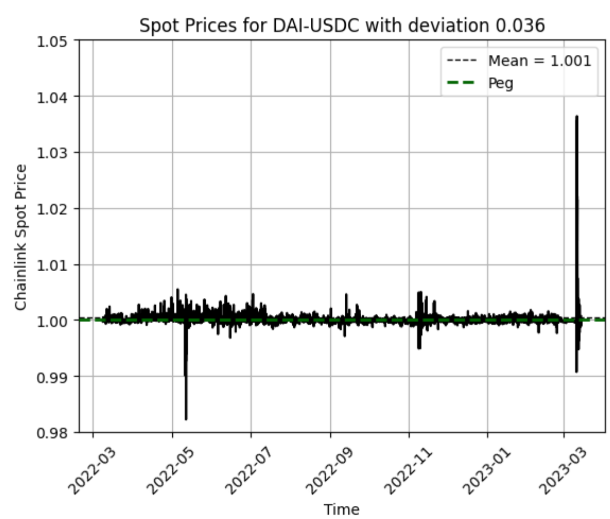

Figure 3: DAI/USDC exchange rate from March 2022 to March 2023 using Chainlink Oracle prices.

2.1 Setting Liquidation Thresholds

Suppose agent A has X_a of token a, and is looking to borrow token b in E-mode. The liquidation threshold for tokens in the E-mode class to which a, b belong is LT. Agent A can then borrow:

$$X_b \cdot p_{b,\text{ETH}} \leq LT \times X_a \cdot p_{a,\text{ETH}} \tag{1}$$

Where $p_{a,\text{ETH}}$ denotes the price of token a in ETH. We know the relative price of tokens a and b is:

$$p_{a,b} = \frac{p_{a,\text{ETH}}}{p_{b,\text{ETH}}} \tag{2}$$

Notice agent A[rb], presumably a statistical arbitrage trader, knows they can avoid the loan-to-value limit by enclosing their borrow within a flashloan transaction. Furthermore, agent A knows that tokens a, b, being part of the same E-mode class, are likely to exhibit strong mean-reverting properties. They wait until there is a large deviation in prices between tokens a and b and then borrow the maximum amount of token b they can with their X_a token a. Soon after, $p_{a,b}$ mean reverts and agent A's collateral gets liquidated. However, they never paid back their X_b loan. Did they profit?

Combining Equations 1 and 2, agent A's borrowing power is:

$$X_b = LT \cdot X_a \cdot (\mu_{a,b} + \gamma_{a,b}) \tag{3}$$

where $\gamma_{a,b}$ is the deviation in price from the long-term mean of $p_{a,b}$, which we denote as $\mu_{a,b}$. Once prices have mean reverted, we can express agent A's profits in terms of token a and find a liquidation threshold to prevent this trade from being profitable:

$$\begin{aligned} X_b \cdot \mu_{b,a} - X_a &< 0 \\ LT \cdot \mu_{b,a}(\mu_{a,b} + \gamma_{a,b}) &< 1 \\ LT &< \frac{1}{1 + \frac{\gamma_{a,b}}{\mu_{a,b}}} \end{aligned} \tag{4}$$

To account for the symmetry in borrowing a for b or b for a, we take the smaller LT:

$$LT = \min\left(\frac{1}{1 + \frac{\gamma_{a,b}}{\mu_{a,b}}}, \frac{1}{1 + \frac{\gamma_{b,a}}{\mu_{b,a}}}\right) \tag{5}$$

Let tokens a, b be DAI, USDC respectively, as in Figure 3, where $\mu_{\text{DAI,USDC}} \approx 1$ and the maximum deviation from March 2022 and March 2023 was $\gamma_{\text{DAI,USDC}} \approx 0.036$. Suppose the community sets a very high liquidation threshold $LT = 0.99$ for stablecoin assets due to their perceived safety. During the collapse of Silicon Valley Bank in March 2023, a trader could have achieved a roughly 2.5% return within the span of a few hours by depositing DAI and borrowing an out-sized amount in USDC. This would leave Aave with roughly 2.5% of whatever capital was deployed as outstanding USDC debt that cannot be repaid with the deposited DAI. For example, \$1M DAI is deposited but \$1.025M USDC can be borrowed and never repaid. Find a snapshot of the oracle prices in Appendix D.

We using Equation 5, and round the LT down to the nearest 0.5% as an additional conservative measure. The maximum liquidation threshold we could offer for an E-mode class containing just USDC and DAI is:

$$LT = \frac{1}{1 + 0.036} \approx 0.9652 \to 0.965 \tag{6}$$

We note that is a risk-off approach, since these deviations are rare, and there are potentially more profitable strategies these traders could execute. Consider two strategies: swap USDT for USDC on Uniswap, or deposit USDT and borrow USDC on Aave. The profitability of either strategy during a momentary USDC de-peg (assuming it will re-peg shortly), depends on the liquidity on Uniswap and the availability for USDC borrows on Aave. Assuming the trader is trading the de-peg in size, there might not be enough sell-side liquidity for USDT on Uniswap, such that the discount the trader suffers on Aave, $1 - LT$, is still more profitable. Our goal in this section was to set an LT such that a trader could not profit on a deviation up to γ.

Notice that by setting the LT according to the largest deviation in prices, we are also increasing the amount of collateral available to liquidate a position when prices are deflated. This minimizes the risk that liquidators will consume all the collateral as prices rise from their local minimum.

2.2 Minimizing Liquidations from a UX/UI Perspective

The LTV is primarily a UX tool to prevent retail traders from taking too much risk, and can be sidestepped by more sophisticated users. Those that borrow at or below the LTV, which from borrowing data we know accounts for many of Aave's borrowers, are presumably not actively managing their margin as much as those borrowing at or near the LT. We can set the LTV such that, based on past extreme deviations, these users still will not get liquidated. We choose extreme deviations as our benchmark as an initial conservative measure, although these deviations are infrequent, and we could achieve greater capital efficiency by basing our methodology off of historical price volatility.

Suppose a user borrows DAI for USDC (instead of agent A borrowing USDC for DAI), when DAI-USDC exchange rates are at or near the long term mean $\mu_{\text{DAI,USDC}} \approx 1$. This user borrows at the LTV ratio, $X_{\text{DAI}} = LTV \cdot X_{\text{USDC}} \cdot \mu_{\text{USDC,DAI}}$. We want to prevent them from getting liquidated. Denote the user's health factor as HF, we set the LTV such that the liquidation $HF = 1$, occurs only at the maximum observed deviation $\gamma_{\text{DAI,USDC}}$:

$$HF = \frac{X_{\text{USDC}} \cdot p_{\text{USDC,ETH}} \cdot LT}{X_{\text{DAI}} \cdot p_{\text{DAI,ETH}}}$$

$$HF = \frac{X_{\text{USDC}} \cdot LT}{LTV \cdot X_{\text{USDC}} \cdot \mu_{\text{USDC,DAI}}} p_{\text{USDC,DAI}}$$

$$1 = \frac{LT}{LTV \cdot \left(1 + \frac{\gamma_{\text{DAI,USDC}}}{\mu_{\text{DAI,USDC}}}\right)}$$

$$LTV \approx \frac{LT}{1 + \gamma_{\text{DAI,USDC}}} \tag{7}$$

$$LTV \approx 0.931 \rightarrow 0.93$$

By setting a liquidation threshold at 0.965 and a LTV at 0.93 we prevent retail users, borrowing at or below their LTV, from getting liquidated at even the most extreme price deviations observed throughout the last year. This is not to say that those borrowing above their LTV will not get liquidated, but they are presumably actively managing their margin and, as we have shown, are unlikely to have an incentive to play adversarial games with Aave. Granted, basing our parameters off the maximum deviation, which occured during the SVB collapse, might be overly conservative. We find that a risk-off approach is a good first step.

2.3 Setting a Liquidation Bonus

The liquidation bonus determines the profitability of liquidators and offsets the potential slippage and gas costs incurred by liquidators. As others have noted, setting the liquidation bonus too high can prevent liquidations from increasing the health factor in the first place, whereas setting them too low can make them unprofitable [3]. Due to the deep on-chain liquidity of stablecoins and LSDs, it is unlikely [as of now] that selling some portion of the aggregated supply of these assets on Aave would incur much slippage.

In this paper, we focus on measuring the mean-reversion of these asset pairs, and ensuring safe LT/LTV parameters. We leave a robust analysis of the 1% LB to future work.

3 Choosing E-mode Assets

As discussed, assets within the same E-mode class must exhibit mean-reversion relative to each other, with an acceptable mean-reversion speed. This mean-reversion gives the protocol additional assurance that (1) the relative value of liens and loans will remain stable over time, and (2) if they diverge, they will eventually converge within an acceptable timeframe.

Adding a new asset to an E-mode class will either reduce its liquidation threshold LT, or leave it unchanged. Although this is not necessarily the case with an alternative framework for setting LT, such as a simulation-based methodology that captures on-chain liquidity projections and borrower behavior, we still need to consider the trade-off of adding new assets to the class.

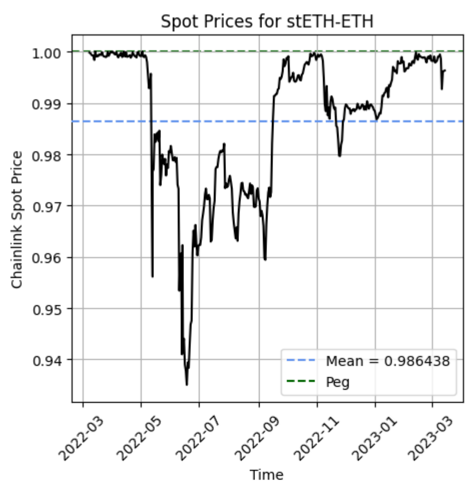

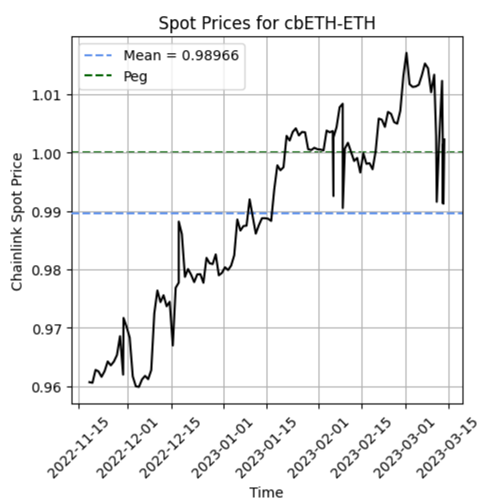

Figure 4: Liquid Staking Derivative prices, for available dates.

3.1 Measuring Mean-Reversion

Mean-reverting times-series are generally characterized as Ornstein-Uhlenbeck (OU) stochastic differential equations:

$$dx_t = \kappa (\mu - x_t) + \sigma dW_t \tag{8}$$

where x_t is an observation of our time-series at time t, κ is the speed of mean-reversion, μ is the long-term mean, σ is a measure of volatility, and W_t is a Weiner process. Intuitively, $\kappa(\mu - x_t)$ is a forcing term that pushes our time-series back towards the mean μ at a "speed" κ, whereas σdW_t introduces some i.i.d. noise term. This equation, or its derivatives such as Vasicek's model or the Cox-Ingersoll-Ross model could be fitted to the price trajectories for E-mode risk simulations. We discuss this as a potential extension in Section 4.

Here, we test for whether a pair of assets would be suitable for E-mode by checking whether their relative prices are mean-reverting, and therefore can be modelled by an Ornstein-Uhlenbeck process. It is not in the scope of this paper to actually produce price projections using O-U (or its derivatives)[1].

There are two widespread methods for measuring mean-reversion in the financial engineering literature. These are the Augmented Dickey-Fuller (ADF) test, and the Hurst exponent [1] [7]. ADF tests for the presence of a unit root in an auto-regressive stochastic process, which in turn determines if said process is stationary, or trend-stationary. A stationary process must also be mean-reverting, making ADF an appropriate test[2]. The Hurst exponent is similarly used to measure mean-reversion, where a Hurst exponent $H < 0.5$ indicates mean-reversion, $H \approx 0.5$ indicates Geometric Brownian Motion (GBM) and $H > 0.5$ indicates a trending process [1].

We implement these in Python using the Statsmodels and Numpy packages. Below are the results from both tests for various stablecoins: USDC, USDT, DAI, LUSD. We are looking for mean-reversion at a 1% confidence interval on ADF as well as a Hurst exponent $H < 0.5$, based on Chainlink oracle prices from March 2022 to March 2023. We re-sample our data at a $45s$ granularity and take a maximum lag on the Hurst exponent of 10000 steps, corresponding to a maximum lag of roughly 100 hours. Since Chainlink oracle prices arrive at irregular timesteps, we discuss our method for resampling our observations in Appendix C.

For liquid staking derivatives, we additionally check that prices must mean-revert relative to ETH, since the underlying asset, WETH, will be included in the E-mode class.

Pair	ADF	p-Value	Hurst	Mean-Reverts
USDC-DAI	-24.660	0.000	0.052	True
USDC-USDT	-10.442	0.000	0.23	True
USDC-LUSD	-4.224	0.001	0.3	True
DAI-USDT	-11.403	0.000	0.214	True
DAI-LUSD	-4.392	0.000	0.3	True
USDT-LUSD	-3.797	0.003	0.284	True

Table 1: Mean reversion results for Stablecoins using Augmented Dickey-Fuller and the Hurst Exponent. The ADF t-score has a 1% critical score of -3.41. Notice that all stables pass the ADF test at a 1% confidence interval, although all LUSD pairs barely pass the ADF test which is reflected in higher Hurst exponents.

Pair	ADF	p-Value	Hurst	Mean-Reverts
stETH-cbETH	-2.352	0.156	0.411	False
stETH-rETH	-2.699	0.074	0.205	False
cbETH-rETH	-2.182	0.213	0.227	False
stETH-ETH	-2.424	0.135	0.46	False
cbETH-ETH	-2.040	0.269	0.368	False
rETH-ETH	-0.253	0.932	0.364	False

Table 2: Mean reversion results for liquid staking derivatives using Augmented Dickey-Fuller and the Hurst Exponent. The ADF t-score has a 1% critical score of -3.41. **That is, there is insufficient evidence of mean-reversion, even for stETH-ETH.**

3.2 Time to Mean-Reversion

Intuitively, mean-reversion is only an attractive property for high efficiency mode assets if the mean-reversion occurs quickly enough. Otherwise, Aave can be left with a hole in its balance sheet for extended periods of time. By measuring the time to mean reversion, we can provide some statistical rigor for why certain tokens do not meet the standards required for E-mode. We can formalize this as follows: for any price deviation $\gamma_{a,b}$ from the mean, we measure the time taken for the price to fall back within a certain tolerance ϵ of the mean $\mu_{a,b}$ within some period τ. We measure this time τ for all deviations

[1] This can be achieved using Maximum Likelihood Estimations, and will likely require the inclusion of a "jump" term to capture the sudden spikes in prices. This jump term is sometimes modelled as a compound Poisson process

[2] A strictly stationary process exhibits both a constant long-term mean, and no drift with respect to time. It follows that each step Δy_t is proportional to the preceding observation y_{t-1}, and tends back towards the mean μ.

in our data. For this paper we have chosen that the $90th$ percentile deviation must converge within $\tau = 24$ hours to a tolerance of 0.1%.

Pair	Avg (hrs)	90th (hrs)	Max (hrs)	Verdict
USDC-DAI	1.056	2.013	18.925	True
USDC-USDT	56.135	158.488	216.062	False
USDC-LUSD	87.870	87.475	2291.250	False
DAI-USDT	4.228	8.764	115.450	True
DAI-LUSD	61.468	49.188	2309.800	False
USDT-LUSD	84.029	84.932	2191.225	False

Table 3: Time to mean-reversion for stablecoins. We measure how long it takes for each pair to revert back to the epsilon-band around its mean, where $\epsilon = 0.1\%$.

We find that only USDC/DAI and DAI/USDT mean-revert within 24 hours at the 90th percentile (as well as on average). Notice that since LSDs don't mean-revert according to our results in Table 2, it does not make sense to test for their mean-reversion speeds.

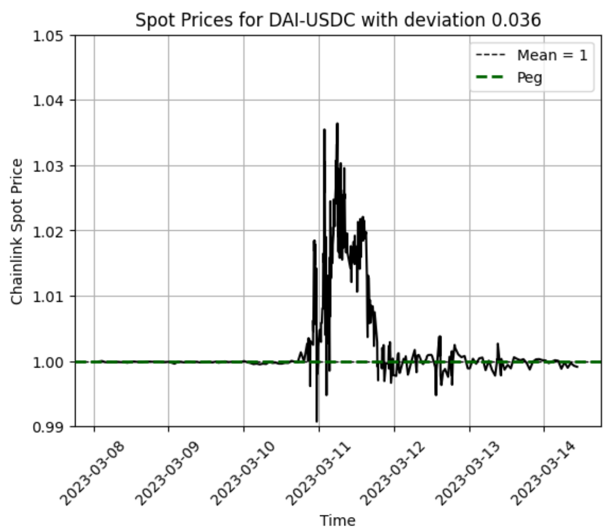

Figure 5: DAI/USDC exchange rate during the March 2023 de-peg. Notice how prices revert to the mean within the matter of hours.

3.3 Defining a Utility Function

We have narrowed our search for E-mode assets to two criteria (not considering deep on-chain liquidity, which is assumed): (1) their relative price with any other asset in the class must be mean-reverting, and (2) they must mean-revert with appropriate speed. Now, we consider a set of eligible assets S, and formulate a utility function to choose which assets should remain in the set, and which can be removed for improved capital efficiency.

Each asset added to the E-mode class would observe an increase in capital efficiency relative to itself, meaning its liquidation threshold and loan-to-value will increase. However, adding this new asset might reduce the increase in capital efficiency for other assets in the E-mode class. Consider the set of stablecoins in an E-mode class, S. Let LT_S be the best liquidation threshold we achieve with this set S, and $LT(s)$ be the "normal-mode" liquidation threshold of asset $s \in S$. A proxy for the appetite to borrow any asset in $s \in S$ is the current collateral supply of that asset. Therefore, the percentage improvement in capital efficiency for each asset $s \in S$ is $\Delta_S LT(s) = LT_S - LT(s)$, whereas the gross improvement to the protocol is $\Delta_S LT(s) \cdot Collateral(s)$[3]. Since it is beneficial to include high-demand assets in E-mode, even if they might lower the overall LT_S, we consider the gross improvement in capital efficiency as our utility function.

$$U(S) = \sum_{s \in S} (LT_S - LT(s)) \cdot \text{collateral}(s) \tag{9}$$

We can define the collateral for some asset A for account i by looking at i's supply of A, denoted as $x_{A,i}$, and their borrowing of other assets $j \in$ assets, denoted as $b_{j,i}$ [4]. We define:

$$c_{A,i} = \frac{x_{A,i} \sum_{j \in \text{assets}} b_{j,i} p_j}{\sum\limits_{j \in \text{assets}} \text{LT}_j \cdot x_{j,i} p_j} \tag{10}$$

Looping through all accounts I, we can glean the overall amount of asset A being used as collateral:

$$c_A = \sum_{i \in I} c_{A,i} \tag{11}$$

Denote the universal set of eligible E-mode assets in a particular asset class, such as stablecoins, as Γ. Our process for choosing E-mode assets based on the proposed utility function is as follows:

1. Consider the power set $\mathcal{P}(\Gamma)$ excluding the empty and unit sets.
2. For each set $S \in \mathcal{P}(\Gamma)$, ensure all pairs $\nu_{a,b} \in \binom{S}{2}$ satisfy our mean reversion criteria.
3. Compute the liquidation thresholds for each pair $\nu_{a,b}$ as per Equation 5. The liquidation threshold for the set, LT_S, is the smallest LT of all pairs in S.

[3] We could additionally multiply each asset's contribution by its reserve factor and current borrow interest rate to measure the expected increase in reserves accrual.

4. Compute the utility generated from LT_S and the collateral data for assets in S.

5. Return the set with the highest utility.

Notice an increase in utility shouldn't come just from assets being used as collateral, as they can also be *borrowed* in E-mode. Including an asset in E-mode creates a significant improvement in UX because users can directly borrow that asset using other E-mode assets as collateral. We are not quantifying this in our utility function. However, notice that assets borrowed in E-mode can then be swapped into the asset that is not in E-mode, achieving a similar increase in capital efficiency - gas and slippage aside. Furthermore, our selection algorithm allows for assets $s \in S'$ to have a lower LT inside the E-mode class than in normal mode. This is convenient for those wishing to borrow this asset, with a higher capital efficiency on their collateral. We do not see why assets in E-mode must enact a higher LT and LTV than outside of E-mode, as specified in the Aave V3 technical specification [2].

3.3.1 Drawbacks

Notice that we might be including new assets in efficiency that cannot currently be used as collateral on Ethereum mainnet, such as LUSD, which will lead to high increases in utility since we set current LT to 0. For USDT, we are using the Avalanche isolation mode liquidation threshold.

Furthermore, we are currently using the supply of assets (as of March 10th, 2023) instead of their collaterals, since the collateral value is not readily available. We leave gathering collateral data for assets for future work. Whether using supply or collateral, this is not a perfect proxy for the actual increase in capital efficiency, since listing assets in E-mode might lead to different impacts in the relative usage of each asset. For example, listing cbETH in E-mode might lead to more relative usage than listing rETH in E-mode, in which case we are understimating the utility of listing cbETH in E-mode.

3.4 Results

Some tokens do not pass our mean reversion tests. A good example LUSD, which does not consistently maintain its peg over time and takes too long to mean-revert, as show in Figure 6.

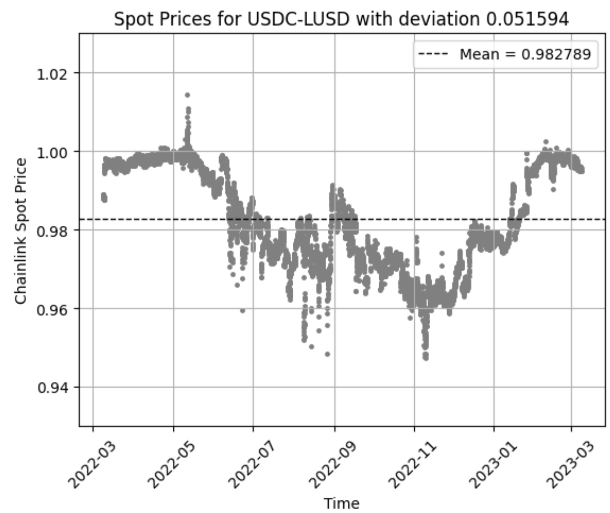

Figure 6: USDC-LUSD exchange rate.

We don't find that any liquid staking derivatives exhibit sufficient mean-reverting properties amongst each other, or relative to ETH. This is likely due to the short history of Chainlink prices on many of these tokens. As price histories increase and token prices stabilize, these LSDs are more likely to pass statistical significance tests such as Dickey-Fuller. On the other hand, we find that all stablecoin pairs pass the ADF test for mean-reversion, and have sufficiently low Hurst exponents. However, only USDC-DAI and DAI-USDT mean-revert sufficiently fast to pass our 24 hour test.

We run our utility function algorithm on the two mean-reverting sets, USDC-DAI and DAI-USDT. Furthermore, we consider potentially including USDC-USDT in our algorithm since it failed our mean-reversion speed test by only a few hours. We obtain the following results:

Pair	LT	LTV	U
USDC, DAI	0.965	0.93	$14,677,250
DAI, USDT	0.87	0.76	$7,550,500
USDC, USDT	0.88	0.775	$2,540,200
USDC, USDT, DAI	0.87	0.76	$8,808,500

Table 4: Results from searching the universe of stablecoins with our utility function.

We conclude that USDC and DAI are the best assets to include in E-mode, with $LT = 0.965$ and $LTV = 0.93$. Both USDC and DAI observe very high deviations with respect to USDT, requiring us to set lower LT, LTV. We argue this tradeoff is not worthwhile. We acknowledge that including USDT could have unquantified benefits in term of convenience for the community, and consider further risks, such as potential de-peg events, in the following section. Further, we acknowledge that using the maximum deviation for setting LT, LTV might also be overly conservative, instead we could use some number of standard deviations, or the $90th$ percentile deviation as we did in Section 3.2.

4 Discussion

Our proposed methodology formalizes the intuition behind E-mode using rigorous statistical tests. However, it does not consider other relevant factors such as liquidity or counterparty risk, and is overly conservative by taking the maximum deviation γ of a pair of assets, instead of perhaps its volatility σ in setting liquidation thresholds and loan-to-value. This becomes especially conservative given the extreme deviations observed due to the Silicon Valley Bank collapse in March 2023, refer to Figure 8.

However, as has been the case with most Aave listings, we find it suitable to start with a risk-off approach before more sophisticated approaches can be built and tested. We find that the Ornstein-Uhlenbeck model discussed in Section 3 is a candidate for improving the diffusion processes baked into existing risk engines. Enforcing mean-reversion properties in exchange rates might produce more accurate price trajectories than pure GARCH models for mean-reverting pairs, as has been shown in some of the algorithmic trading literature [12]. Below we consider a few potential concerns and improvements.

4.1 Fat Tail Events

Fat tail events, sometimes referred to as black swan events, are highly improbable events that carry with them significant risks. For Aave, this risk refers to the risk of a catastrophic asset devaluation event, such as a UST-like de-peg for one of Aave's collateral assets. This risk is a function of a token's market cap, liquidity, centralization, or fear regarding the issuing party's reserves (e.g. Tether). These have not been accounted for in our methodology. We have contained our research to the highest market-cap stablecoins and LSDs to mitigate potential counterparty risk, as described in Section 1.1.

We argue the debt ceiling in Aave's isolation mode is the better line of defense against fat tail risks. Placing USDT both in high efficiency mode and isolation mode, for example, allows users to benefit from the increased capital efficiency of borrowing/lending a mean-reverting pair, while mitigating the fat tail risk of a potential USDT de-peg. For this reason, we leave managing fat tail risks to further work on Aave's isolation mode. Other avenues for fat tail risk management are potential examinations of collateral assets and their associated counterparty risks, including reviews of reserve audits and/or regulatory compliance.

4.2 Future Work

Setting the minimum speed of mean reversion is a relevant lever for E-mode. Demanding E-mode assets mean revert at faster rates would reduce the probability that Aave is left holding bad debt for longer, but also reduces the potential assets we would include in an E-mode class. We have chosen a period of 24 hours arbitrarily, and are open to more sophisticated ways of setting this parameter. Furthermore, we have considered the pairwise exchange rates of assets and their pairwise mean-reversion. However, these E-mode classes could include > 2 assets. A better approach could consider portfolios that include more than just two pairs, sampled from some appropriate distribution.

As previously mentioned, integrating E-mode into the agent based simulations of Aave's risk providers will likely yield more appropriate parameters than what has been proposed in Section 2. Liquidation thresholds are set such that they minimize the protocol's value at risk, given each asset's historical on-chain liquidity and asset price volatility. We find that taking stressed VaR measurements from agent based simulations, as previously discussed in [5], would be a more sophisticated approach.

4.2.1 E-mode as an Ornstein-Uhlenbeck Process

The Cox–Ingersoll–Ross (CIR) model is an extension of the Ornstein-Uhlenbeck process and belongs to a subset of jump diffusion processes. Since diffusion processes like GARCH or jump diffusion are what Chaos Labs is currently building to model risk parameters for Aave, we find this to be an attractive potential extension to the simulations when it comes to E-mode.

We can use statistical methods to estimate the parameters for these models, such as mean-reversion speed κ or instantaneous volatility σ, and project mean-reverting pair prices. Common methods found in the literature for making these estimates include ordinary least squares regression [10] or maximum likelihood estimators [10] with various probability densities. This could potentially lead to extensions to other mean-reverting pairs, such as WBTC/WETH, which would have $\mu \neq 1$ and might exhibit regime shifts over time (i.e. the mean isn't always constant).

4.2.2 Adding Jumps

We modelled our process as a simple mean-reverting series in equation 8. This doesn't capture the sudden rare spikes in prices which we usually term as de-pegs. Adding a jump term to our OU process is one option to capture these jumps in simulations:

$$dx_t = \kappa (\mu - x_t) + \sigma dW_t + dJ_t \tag{12}$$

Where J is sometimes modelled as a compound Poisson process.

A E-Mode Class Algorithm

Algorithm 1: **E-Mode Class Optimization**

$\Gamma \leftarrow \{\text{Universe of eligible assets}\}$
$S^* \leftarrow \{\}$ ▷ Current best set
$U^* \leftarrow 0$ ▷ Current best utility
$Q \leftarrow \{\}$ ▷ Set of ineligible pairs
for $S \in \mathcal{P}(\Gamma)$ **do**
$LT_S \leftarrow 1$
$LTV_S \leftarrow 1$
success $\leftarrow$ **true**
for $\nu_{s_i,s_j} \in \binom{S}{2}$ **do**
▷ Tuple ν_{s_i,s_j} is a pair in S.
if $\nu_{s_i,s_j} \in Q$ **then**
▷ We know ν_{s_i,s_j} fails our criteria.
success $\leftarrow$ **false**
break
end if
$\mathbb{1}(\nu_{s_i,s_j})$ ▷ Mean-Reversion results
if $\mathbb{1}(\nu_{s_i,s_j}) =$ **true then**
$LT \leftarrow$ LT from eq 5
$LTV \leftarrow$ LTV from eq 7
$LT_S \leftarrow \min(LT_S, LT)$
$LTV_S \leftarrow \min(LTV_S, LTV)$
else
$Q \leftarrow Q \cup \{\nu_{s_i,s_j}\}$
success $\leftarrow$ **false**
break
end if
end for
$U \leftarrow \sum_{s \in S} (LT_S - LT(s)) \cdot Collateral(s)$
if success = **true and** $U > U^*$ **then**
$U^* \leftarrow U$
$S^* \leftarrow S$
end if
end for
return S^*

Since the power set of stablecoins or LSDs is small, and the algorithm is relatively computationally inexpensive, optimizing it is left for future work.

B Tether Holdings Limited Reserves Audit

Asset Category	Amount in USD
1. Cash & Cash Equivalent & Other Short-Term Deposits & Commercial Paper	
U.S. Treasury Bills[4]	39,230,259,046
Commercial Paper and Certificates of Deposit[5]	-
Money Market Funds[6]	7,372,926,391
Cash & Bank Deposits[7]	5,318,311,794
Reverse Repurchase Agreements[8]	3,046,093,954
Non-U.S. Treasury Bills[9]	93,849,833
Subtotal	**55,061,441,018**
2. Corporate Bonds, Funds & Precious Metals	3,444,097,599
3. Other Investments	2,685,786,230
4. Secured Loans	5,852,823,328
Total (1+2+3+4)	**67,044,148,175**

Figure 7: Reserve assets of Tether Holdings Ltd. by Binder Dijker Otte (BDO), on December 2022.

C Regularizing Chainlink Oracle Prices

Chainlink oracle prices arrive at irregular times. Their intervals are defined as heartbeats. For example, every 24 hours there is a heartbeat for the USDC-USD oracles, at which point oracles must update their prices on chain. This heartbeat is 1 hour for DAI-USD. To get our USDC-DAI exchange rate, we forward fill USDC-USD prices at the same interval as DAI-USD. However, oracles are also required to update their prices whenever prices deviate above some predefined threshold, 0.25% for USDC and DAI. During period of volatility, oracles update very frequently.

When conducting statistical tests such as Dickey-Fuller, or when measuring speed of mean reversion using OLS, we are assuming regular time intervals. We resample our data at a 45 second granularity, and use this regularized timeseries for our statistical testing. 45 seconds was chosen as it retains information on the moments of peak volatility without being too computationally demanding.

D USDC-DAI Momentary De-Peg Data

datetime	DAI	USDC	USDC-DAI	Deviations
2023-03-11 05:09:23	0.9287	0.9099	0.9798	0.0197
2023-03-11 05:12:11	0.9316	0.9099	0.9767	0.0228
2023-03-11 05:28:47	0.9340	0.9095	0.9737	0.0257
2023-03-11 05:32:47	0.9374	0.9095	0.9702	0.0293
2023-03-11 05:34:47	0.9314	0.9095	0.9764	0.0230
2023-03-11 05:36:23	0.9353	0.9095	0.9724	0.0271
2023-03-11 05:44:59	0.9397	0.9095	0.9679	0.0316
2023-03-11 05:45:47	0.9372	0.9095	0.9704	0.0291
2023-03-11 05:57:23	**0.9398**	**0.9068**	**0.9649**	**0.0346**
2023-03-11 05:57:47	0.9358	0.9068	0.9691	0.0304
2023-03-11 05:59:47	0.9324	0.9068	0.9726	0.0269
2023-03-11 06:01:47	0.9352	0.9084	0.9714	0.0281
2023-03-11 06:02:47	0.9306	0.9084	0.9761	0.0233
2023-03-11 06:03:23	0.9352	0.9084	0.9714	0.0281
2023-03-11 06:06:23	0.9324	0.9084	0.9742	0.0252
2023-03-11 06:08:23	0.9297	0.9060	0.9745	0.0250
2023-03-11 06:09:23	0.9326	0.9060	0.9714	0.0280
2023-03-11 06:10:23	0.9277	0.9060	0.9766	0.0229
2023-03-11 06:11:47	0.9306	0.9102	0.9780	0.0215
2023-03-11 06:13:23	0.9276	0.9102	0.9812	0.0183
2023-03-11 06:14:47	0.9338	0.9081	0.9724	0.0270
2023-03-11 06:15:23	0.9313	0.9081	0.9750	0.0245
2023-03-11 06:16:47	0.9343	0.9081	0.9720	0.0275
2023-03-11 06:17:23	0.9304	0.9081	0.9760	0.0235
2023-03-11 06:17:47	0.9279	0.9081	0.9787	0.0208
2023-03-11 06:18:47	0.9233	0.9081	0.9835	0.0159
2023-03-11 06:19:23	0.9300	0.9081	0.9765	0.0230
2023-03-11 06:21:23	0.9343	0.9081	0.9719	0.0276
2023-03-11 06:29:47	0.9254	0.9071	0.9803	0.0192

E Stablecoin Exchange Rates

Notice the prices during the SVB collapse on March 2023:

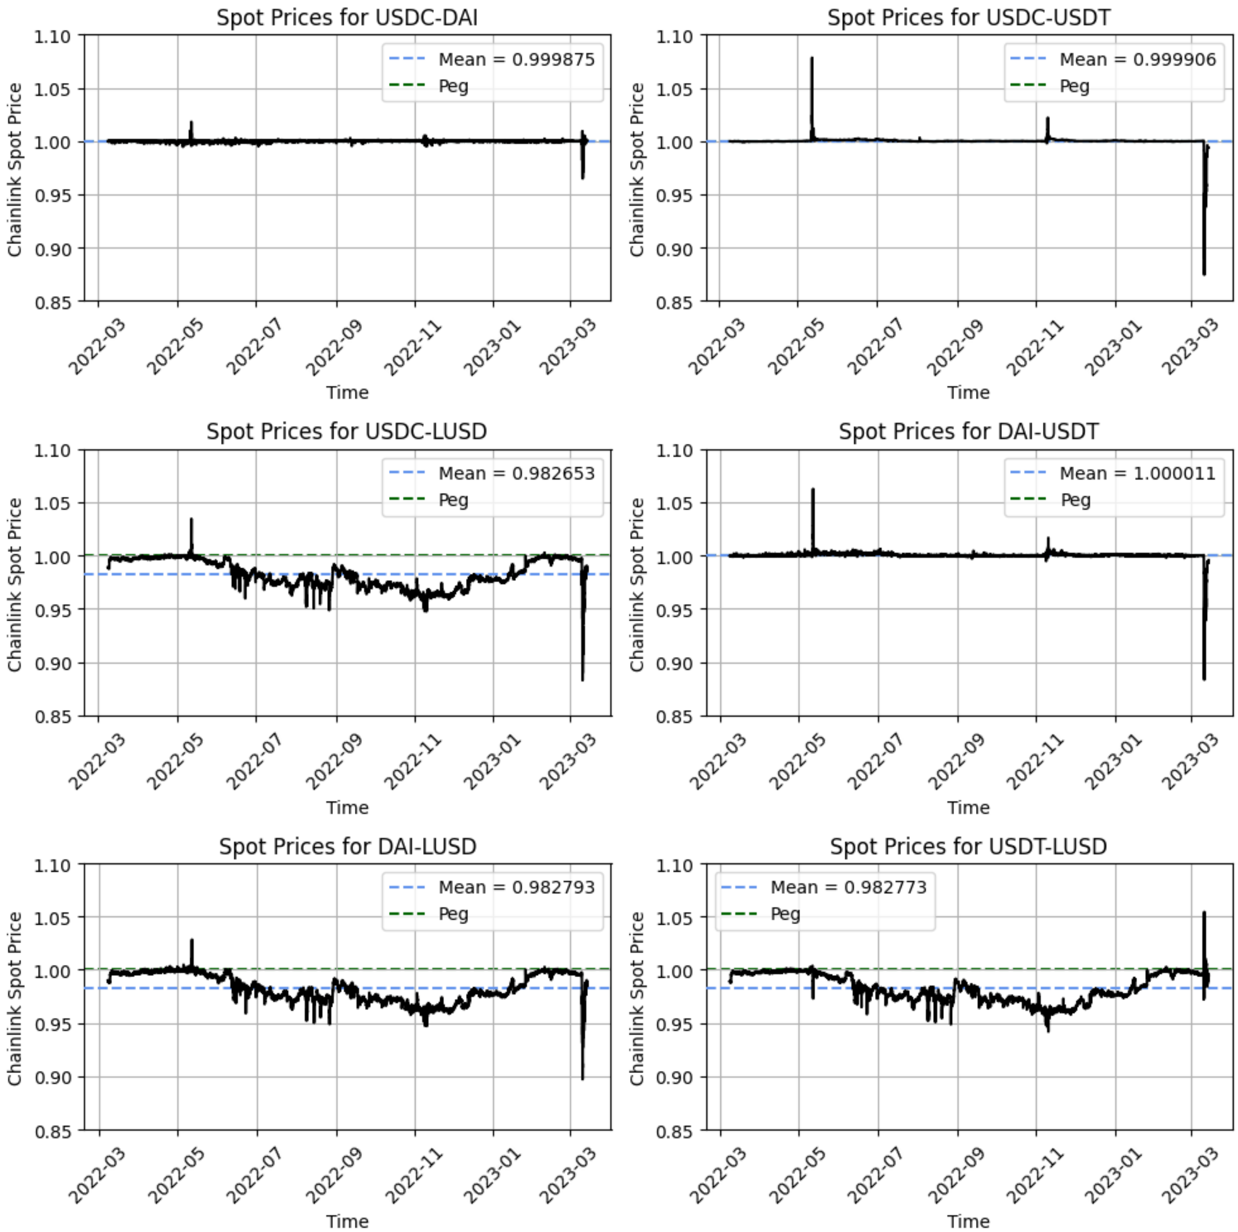

Figure 8: Stablecoin exchange rates throughout March 2022 to March 2023.

F Authors

In alphabetical order.

Omer Goldberg omer@chaoslabs.xyz

Yonatan Haimowitz haimo.yonatan@chaoslabs.xyz

References

[1] Ernest Chan. *Algorithmic Trading*. New Jersey: John Wiler and Sons, 2013.

[2] Emilio Frangella and Lasse Herskind. “Aave V3 Technical Paper”. In: (2022).

[3] Gauntlet. “Market Risk Assessment”. In: (2021).

[4] Yonatan Haimowitz and Omer Goldberg. “A Note on the Aave Supply Cap Parameter”. In: (2023).

[5] Yonatan Haimowitz and Omer Goldberg. “Aave V3 Risk Parameter Methodology”. In: (2023).

[6] P. Koning J. *What Tether means when it says it is Regulated*. URL: https://www.coindesk.com/policy/2021/01/26/what-tether-means-when-it-says-its-regulated/. accessed:03/12/2023.

[7] Kris Longmore. *Exploring mean reversion and cointegration with Zorro and R: part 1*. URL: https://robotwealth.com/exploring-mean-reversion-and-cointegration-with-zorro-and-r-part-1/. accessed:03/09/2023.

[8] Shaurya Maulwa. *Tether’s USDT Stablecoin Drops 3% Below $1 Peg*. URL: https://www.coindesk.com/markets/2022/11/10/tethers-usdt-stablecoin-slips-from-1-peg/. accessed:03/12/2023.

[9] John Morrow. *[ARFC] Create eMode Categories for Aave v3 on Ethereum*. URL: https://governance.aave.com/t/arfc-create-emode-categories-for-aave-v3-on-ethereum/10981/3. accessed: 03/08/2023.

[10] Kostas Triantafyllopoulos and Giovanni Montana. *Dynamic modeling of mean-reverting spreads for statistical arbitrage*. 2008. DOI: 10.48550/ARXIV.0808.1710. URL: https://arxiv.org/abs/0808.1710.

[11] 117th Congress of the United States. *Stablecoin TRUST Act of 2022*. URL: https://www.banking.senate.gov/imo/media/doc/the_stablecoin_trust_act.pdf. accessed:03/12/2023.

[12] Jianbo Wang and Jianyang Fang. *The application of Ornstein-Uhlenbeck Process model and ARCH/GARCH model in statistical arbitrage*. 2017. URL: https://lup.lub.lu.se/luur/download?func=downloadFile&recordOId=8913459&fileOId=8913462.