

A Note on Isolation Mode Risk Parameters

Chaos Labs

February 2023

Abstract

The Aave V3 protocol incorporates an innovative feature known as isolation mode, which allows the inclusion of high-risk assets while implementing safeguards to curtail potential losses incurred by the protocol. This mode serves as a risk mitigation measure by instituting a debt ceiling, delineating the upper limit of borrowable value associated with a given asset. Determining an asset's debt ceiling necessitates striking a delicate balance: it should facilitate meaningful borrowing activities while mitigating the risk of exposing the protocol to an untenable level of losses in the event of a catastrophic occurrence. This document presents a prudent framework that the Aave protocol can adopt to systematically establish appropriate debt ceilings for its diverse range of assets.

1 Introduction

Aave V3 Protocol introduces a noteworthy addition known as *isolation mode* [FH22] within the Aave ecosystem. This novel functionality offers a distinct avenue for asset listing on the Aave protocol, wherein assets designated for isolation mode can be supplied in a customary manner. Yet, their utilization is restricted solely to collateralizing stable assets. As isolation mode assets inherently carry a higher degree of risk than their non-isolation mode counterparts, the protocol incorporates a risk parameter referred to as the *debt ceiling*. This parameter is a crucial mechanism, establishing an upper threshold denominated in USD value, limiting the borrowing capacity against each isolation mode asset. To illustrate, consider the scenario where asset A is enlisted under isolation mode with a debt ceiling set at \$15 million. Consequently, asset A can only be utilized to collateralize a cumulative amount of stablecoins equating to \$15 million across all borrowers active on a given Aave deployment. Please refer to Figure 1 for a visual representation of the borrowing capabilities associated with isolated assets.

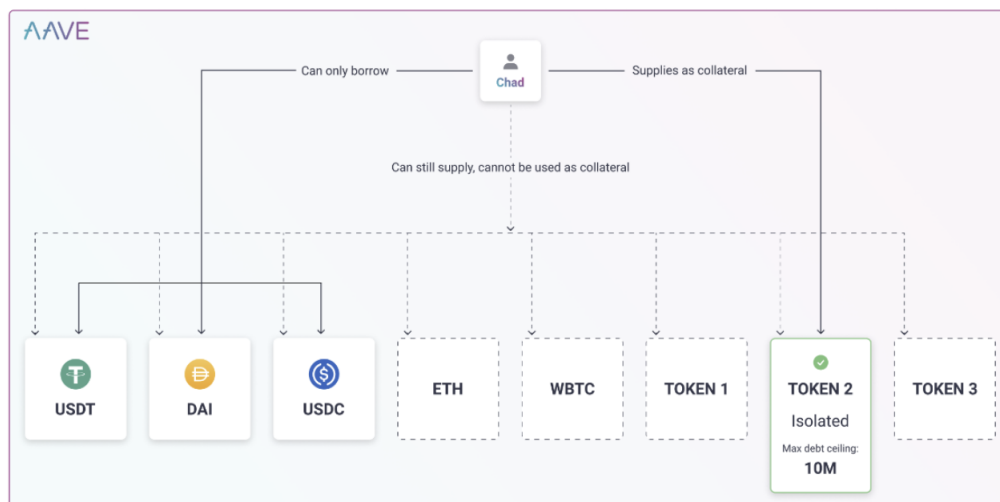


Figure 1: An illustration of an isolated asset TOKEN 2 on the Aave V3 Protocol [FH22].

In the context of the Aave V3 Protocol, isolation mode assets refer to those assets that are identified by the protocol as having an inherent susceptibility to a highly adverse event resulting in a *catastrophic devaluation*. These assets, due to their elevated risk profile, are subject to a stringent risk management framework implemented by the protocol. Consequently, the debt ceiling emerges as a pivotal parameter, representing the uppermost limit denoting the potential loss that may be incurred by the protocol and its lenders in the unfortunate occurrence of such a catastrophic asset devaluation event. The debt ceiling thus serves as a safeguard, enabling the protocol to control and mitigate its exposure to risk prudently.

Definition 1 (Catastrophic Asset Devaluation Event). *A catastrophic asset devaluation event refers to a critical situation characterized by a swift and substantial decline in the price of an asset. Such events manifest when assets experience a rapid deterioration in value, leading to adverse implications for the overall market and participants involved. Notable instances of catastrophic asset devaluation include the collapse of tokens such as Terra UST and LUNA, where the prices of these assets experienced a precipitous decline. In these particular scenarios, the devaluation of the asset was so severe that it resulted in a significant reduction in both price and liquidity. Consequently, the liquidators operating within the Aave protocol encountered significant challenges in profitably liquidating borrow positions collateralized by the devalued asset. The inability to effectively execute these liquidations contributed to extensive losses experienced across the protocol, affecting multiple participants. These instances exemplify the potential risks associated with catastrophic asset devaluation events and underscore the significance of implementing robust risk management to mitigate widespread protocol losses.*

While isolation mode assets can indeed be borrowed, it is important to acknowledge that a rapid increase

in the valuation of such assets also carries inherent risks, albeit accompanied by potential additional revenue in the form of borrow fees. To effectively manage these risks, the Aave protocol incorporates borrow and supply caps as precautionary measures. These caps constrain the borrowing and lending activities associated with isolation mode assets, ensuring that the protocol can prudently handle sudden valuation surges while maximizing revenue generation.

Furthermore, it is noteworthy that revenue generated from borrowing isolated assets, in terms of interest accrued to the treasury collector, remains unaffected regardless of whether these assets can be utilized as collateral. This aspect underscores that revenue generation is not contingent on the collateral usability of isolated assets, thereby providing a consistent income for the protocol.

By leveraging isolation mode in conjunction with implementing the debt ceiling, the Aave protocol effectively establishes bounds on potential losses for assets exhibiting a non-negligible probability of experiencing catastrophic asset devaluation events. These risk management mechanisms play a crucial role in safeguarding the protocol and mitigating the extent of potential losses that could be incurred.

Among the various risk parameters governing the Aave protocol, the debt ceiling stands as a distinctive feature exclusive to isolation mode. It serves as a dedicated mechanism designed to mitigate the protocol's exposure specifically to catastrophic asset devaluation events (Aave Risk Methodology, date). In addition to the debt ceiling, the protocol employs other risk parameters to manage its exposure to different economic risks effectively. For instance, the liquidation threshold assumes the role of controlling the protocol's market and price exposure. In contrast, supply caps and the liquidation bonus contribute to managing liquidity risk (Aave Risk Methodology [Comb]).

It is important to note that the debt ceiling's primary purpose lies in containing the protocol's potential exposure to catastrophic asset devaluation events, distinguishing it from other risk parameters. Methodologies have already been implemented for determining the liquidation threshold and liquidation bonus for all assets integrated into Aave V3 [HG23]. In this context, we present a comprehensive methodology to evaluate the fat tail risk associated with isolation mode assets, providing valuable insights to inform the establishment of appropriate debt ceilings. The assessment framework is specifically tailored to limit the potential fallout resulting from catastrophic asset devaluation events, thereby enhancing the protocol's resilience and risk management capabilities.

1.1 The Current State of Isolation Mode

The deployment of isolation mode is currently focused on a limited selection of assets within the Aave v3 ecosystem, specifically targeting markets on Avalanche, Arbitrum, Polygon, and Optimism (Aave Docs, date). It is important to note that the provided pool configurations may be subject to potential updates and modifications. Thus the information presented in the Aave documentation might not reflect the most recent developments [Coma].

To provide a benchmark, it is worth mentioning that the existing debt ceiling for USDT on the Avalanche network has already been reached, amounting to \$5 million. This observation has prompted community discussions and engagement to explore potential avenues for raising the debt ceiling limit in order to accommodate the growing demand and requirements of the platform. The collaborative efforts from the community indicate an active involvement in shaping the protocol's governance and risk management policies.

2 Setting the Debt Ceiling

The debt ceiling plays a pivotal role in mitigating market risks associated with the inclusion of risky collateral within the Aave protocol. Determining an appropriate debt ceiling necessitates a comprehensive understanding and quantification of the benefits of listing these assets and the inherent risks associated with sudden devaluation events. However, quantifying the probabilities of catastrophic asset devaluation events presents considerable challenges, thereby requiring alternative approaches.

In this regard, we propose adopting a stress testing framework, specifically a sensitivity analysis, to assess the potential losses incurred by the protocol in the event of a catastrophic devaluation. Additionally, we introduce potential constraints aimed at mitigating excessive risk exposure. Our framework is grounded in conservative liquidity assumptions, drawing insights from the LUNA token crash that transpired on May

9th, 2022. By leveraging these conservative liquidity assumptions, we aim to provide the community with a range of boundaries delineating the possible protocol losses, often referred to as bad debt, in situations where the outstanding debt ceiling must be liquidated under stressful conditions. This endeavor aims to offer valuable insights to stakeholders, enabling them to comprehend the potential ramifications of adverse scenarios better and effectively manage associated risks.¹

2.1 Problem Setup

We begin with some definitions and assumptions.

d	The debt ceiling parameter, with units of US dollars.
T	Some pre-defined time interval; we set $T = 1$ year for the purposes of our analysis.
p	The probability of a catastrophic asset devaluation event for the risky asset in question at a time $\tau \in (0, T)$.
τ	A random variable representing the time at which the catastrophic asset devaluation event occurs, where $\tau = 0$ would mean that a catastrophic event occurs immediately after listing the asset in isolation mode, and $\tau = T$ means that a catastrophic event occurs T years after listing the asset in isolation mode.
r_d	The reserve collection per unit of time for the debt ceiling d , in US dollars. This accrues to the pool's treasury collector, which can roughly be seen as revenue.
$L(d)$	The proportion of the debt ceiling the protocol would lose due to a catastrophic event; $L \in [0, 1]$.

Table 1: Variable Definitions

These variables are subject to the following assumptions:

1. **Constant Revenue:** The revenue per unit of time, r_d , is constant over time. We perform a sensitivity analysis on this assumption using historical USDC borrow rates across Aave markets in section 2.3.
2. **Uniform Distribution:** τ is uniformly distributed over $(0, T)$.
3. **Atomic Liquidation:** When a catastrophic asset devaluation event occurs, the protocol loses a large portion of the debt ceiling, parameterized by $L(d)$. We estimate a conservative loss curve $L(d)$ for the case where LUNA was provided as isolated collateral during its crash in May 2022. Refer to section 2.2 for our methodology.

Given these definitions and assumptions, let us see the protocol's expected reserves accrued at time T for listing the risky asset with a debt ceiling d . These reserves are accrued by the reserve factor to the treasury collector contract of each pool.

$$\begin{aligned}
 \mathbb{E}[\text{reserves at time } T] &= p(\mathbb{E}[r_d\tau - dL(d) \mid \text{catastrophic event occurred}]) + (1 - p)r_dT \\
 &= p(r_d \cdot \mathbb{E}[\tau \mid \text{catastrophic event occurred}] - dL(d)) + (1 - p)r_dT \\
 &= p(r_d \cdot (T/2) - dL(d)) + (1 - p)r_dT \\
 &= pr_dT/2 - pdL(d) + r_dT - pr_dT \\
 &= r_dT - p(r_dT/2 + dL(d)).
 \end{aligned}$$

We have constructed an expectation for the reserves accrued to the protocol from the risky asset as a function of a loss parameter $dL(d)$ ². We set a conservative upper bound on the debt ceiling such that

¹Notice our framework is meant to supplement the existing risk methodologies employed by Aave, Chaos Labs, and Gauntlet, specifically for assessing fat tail events on risky assets.

²Technically, $\mathbb{E}[\tau \mid \text{catastrophic event occurred}]$ is slightly less than $T/2$, due to the [Lindy Effect](#). However, we expect that accounting for this effect would not yield a meaningful impact on the zero-profit debt ceiling.

the expected reserves accrued are greater than or equal to zero given the probability of a catastrophic devaluation³. That is, $r_d T - p(r_d T/2 + dL(d)) \geq 0$, and so $r_d T - pr_d T/2 \geq pdL(d)$, and thus:

$$d \leq \frac{r_d T(1/p - 1/2)}{L(d)}. \quad (1)$$

Denote the point at which our expected reserves are 0 as the breakeven point. We propose that the breakeven point is a useful metric in guiding how to set our debt ceiling: it is a function of the asset’s perceived fat tail riskiness, p , and the adverse effect of its losses, $L(d)$.

$$d = \frac{r_d T(1/p - 1/2)}{L}. \quad (2)$$

2.2 Estimating a Loss Parameter

Suppose Aave had listed LUNA as an isolation mode asset in early 2022. In May 2022, an unprecedented event occurred where the price of LUNA experienced a drastic crash, reaching a value of 0 within a matter of days. As a result, the outstanding debt associated with LUNA collateral needed to be repaid promptly during this period. To adopt a highly conservative stance, we can assume that at some point during LUNA’s crash, the entire debt ceiling has defaulted, requiring its liquidation in a single Ethereum transaction.

In light of this scenario, we aim to model the gross amount of the borrowed asset that a profit-maximizing liquidator would be willing to repay under the assumption that the borrowed asset is a liquid stablecoin, such as USDC. By considering the perspective of a profit-maximizing liquidator, we can evaluate the potential amount they would be willing to offer for repayment in order to optimize their own gains within the given market conditions. This analysis enables us to gain insights into the possible outcomes and implications of a liquidation event associated with LUNA’s crash, facilitating a comprehensive understanding of the potential risks and considerations involved.

The most liquid venue for swapping LUNA was a LUNA/WETH Uni v3 pool⁴. We model the liquidation transaction as a “flash swap” on Uniswap.

Flash Liquidation: A LUNA liquidation using a LUNA/WETH Uni v3 pool.

1. The liquidator initiates a flash swap by borrowing η_{LUNA} LUNA.
2. They swap η_{LUNA} LUNA for η_{ETH} ETH.
3. They swap the η_{ETH} ETH for η_{USDC} USDC.
4. They repay η_{USDC} to the protocol and receive $(1 + LB)\eta_{USDC} \cdot p_{LUNA/USDC}$, where LB is the liquidation bonus and $p_{LUNA/USDC}$ is the amount of LUNA they receive for 1 USDC (an oracle price).
5. They repay their debt to Uniswap, for a profit of $(1 + LB)\eta_{USDC} \cdot p_{LUNA/USDC} - \eta_{LUNA}$.

To understand the profit-maximizing liquidator, we overview Uniswap v3 mechanics in appendix A. At a high level, the profit-maximizing liquidator performs a convex optimization on their profit curve under the constraints of the underlying AMM architecture. The output of this optimization is the amount of USDC the liquidator will repay to the protocol. We plot our results in 2.

2.2.1 Adjusting for Relative Liquidity

In our analysis thus far, we have made the assumption that liquidations align with the liquidity profile observed in the Uni v3 pool for the WETH/LUNC pair during its crash in May 2022. This approach has proven valuable in capturing the liquidity dynamics and asymmetry exhibited by LUNA during that specific

³We discuss this assumption later, factoring in the more subjective benefits of listing isolation mode assets, such as growing the protocol’s ecosystem.

⁴Contract address: 0x16b70f44719b227278a2dc1122e8106cc929ecd1

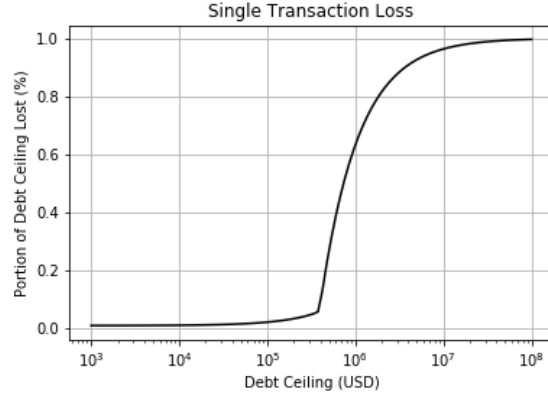


Figure 2: Bad debt remaining after a single profit-maximizing liquidation transaction on a Uni v3 LUNA/WETH pool at block 14740860. This assumes a liquidation threshold of 80% and a liquidation bonus of 15%.

period compared to its liquidity levels in the preceding months. However, it is important to acknowledge that larger tokens, such as UNI, typically possess significantly higher liquidity levels. In appendix A we derived the optimal amount of token1 we get by crossing each tick on a Uni v3 pool. We can rearrange it to show that the optimal amount of token0 we get out of the pool scales linearly with the liquidity at each tick:

$$\delta'_{x,i} = \min \left[\lambda_i \cdot \left(\frac{1}{\sqrt{p_i}} - \frac{1}{\sqrt{\alpha(LB+1)}} \right), \lambda_i \cdot \left(\frac{1}{\sqrt{p_i}} - \frac{1}{\sqrt{p_{i+1}} - \sqrt{p_i}} \right) \right] \quad (3)$$

To account for the liquidity differences between various tokens and adjust our estimate for the loss parameter under severely stressed conditions, we can scale the liquidity in each bucket by a liquidity ratio. This ratio allows us to adjust the profitable repayments based on the liquidity of different tokens.

One approach to determining the liquidity ratio is by considering measures such as the 2% depth or Total Value Locked (TVL) held on major liquidity sources for each token. Since TVL is readily accessible through querying, it can serve as a useful metric for generating liquidity ratios and adjusting the loss parameter for each token.

For instance, we can analyze the TVL of LUNA in the early months of 2022 before its crash and compare it to the TVL of tokens like UNI, MKR, and others during February 2022. If UNI currently exhibits approximately double the on-chain TVL compared to LUNA’s pre-crash TVL, we can approximate scaling LUNA’s mid-crash liquidity by a factor of approximately 2 to obtain a more accurate estimate of the loss parameter for UNI.

This approach leverages available data on TVL to establish relative liquidity ratios, offering a more nuanced understanding of the liquidity dynamics across different tokens. By adjusting the loss parameter accordingly, we can better capture the potential outcomes and considerations associated with severely stressed conditions, providing valuable insights for risk assessment and management within the Aave protocol.

We find that UNI had $\approx 10\times$ higher TVL than LUNA during the stated intervals. This is a slight underestimate of the actual TVL ratio since we only considered the top 4 liquidity pools for UNI on Uniswap, but these accounted for the vast majority of TVL. For each token, we consider its adjusted loss curve instead of naively extrapolating from LUNA’s.

2.2.2 Considerations

Our current formulation exhibits a high level of conservatism. In practical scenarios, it is important to consider the existence of multiple liquidity sources and the involvement of diverse liquidators capable of conducting position liquidations over multiple blocks. Notably, a significant increase in the ETH/LUNA

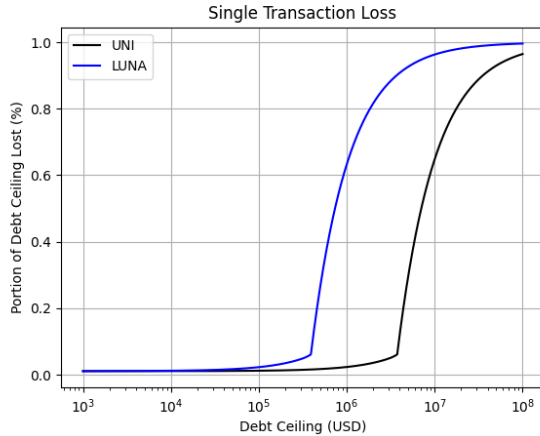


Figure 3: Bad debt remaining after a single profit-maximizing liquidation transaction for UNI vs LUNA, where UNI liquidity is scaled by a liquidity ratio derived from on-chain TVL.

swap price (leading to a corresponding decrease in the LUNA/ETH swap price) would create a favorable arbitrage opportunity for traders. These traders could exploit this situation by swapping their ETH for LUNA within the pool, effectively acquiring LUNA at a discounted rate, and subsequently selling it at a profit on alternative trading platforms, such as centralized exchanges. This process of re-balancing implies that liquidations executed across multiple transactions and blocks could potentially result in lower losses compared to those depicted in Figure 2.

Additionally, it is overly conservative to assume that borrowers would not rebalance their positions to evade liquidation. While instances of adversarial games, like the CRV attack observed in late 2022, demonstrate deliberate avoidance of liquidation by borrowers, in general, borrowers tend to take precautions to avoid liquidation due to the haircut incurred from the liquidation bonus.

To address these considerations, we outline several potential enhancements to our methodology in section 3. Nevertheless, despite its limitations, our current approach serves as a useful upper bound for estimating potential protocol losses. This upper bound can effectively guide discussions surrounding debt ceiling decisions until a more sophisticated approach is proposed. In subsequent sections, we will employ this loss curve as an upper bound to facilitate risk assessments, providing valuable insights into risk evaluation within the protocol.

2.3 Reserve Trajectories

The reserve factor on the borrowed asset dictates the percentage of the borrow interest rate allocated to the collector contract, which serves as a reserve for the protocol. For relevant stablecoins (USDC, DAI), the reserve factor is set at 10% [Coma]. It follows that our parameter r_d can be expressed as a function of the borrowing rate r and debt ceiling d :

$$r_d = 0.1rd \tag{4}$$

We can query historical borrow rates from the Aave reserve API. In figure 4, we show that stablecoin rates have been approximately bounded in the range $[0.5\%, 2\%]$ over the last 6 months. The simple average stablecoin rate for USDC and DAI over Avalanche, Ethereum, and Polygon is $\approx 1\%$.

As a benchmark, the annualized reserves collected from listing USDT as an isolated asset on Avalanche with a debt ceiling of $\$5M$ is likely to fall somewhere in $r_d \in [\$2, 500, \$10, 000]^5$. Notice that some borrowers opt for stable rate borrows which mostly sit at $\approx 5\%$.

⁵As a sanity check, we can show that Aave’s outstanding borrowing volume as of February 27th, 2022 is $\$2.19B$. Assuming a roughly 3% borrow rate (on the upper end since non-stablecoins usually have higher rates) and reserve factor of 15% (again,

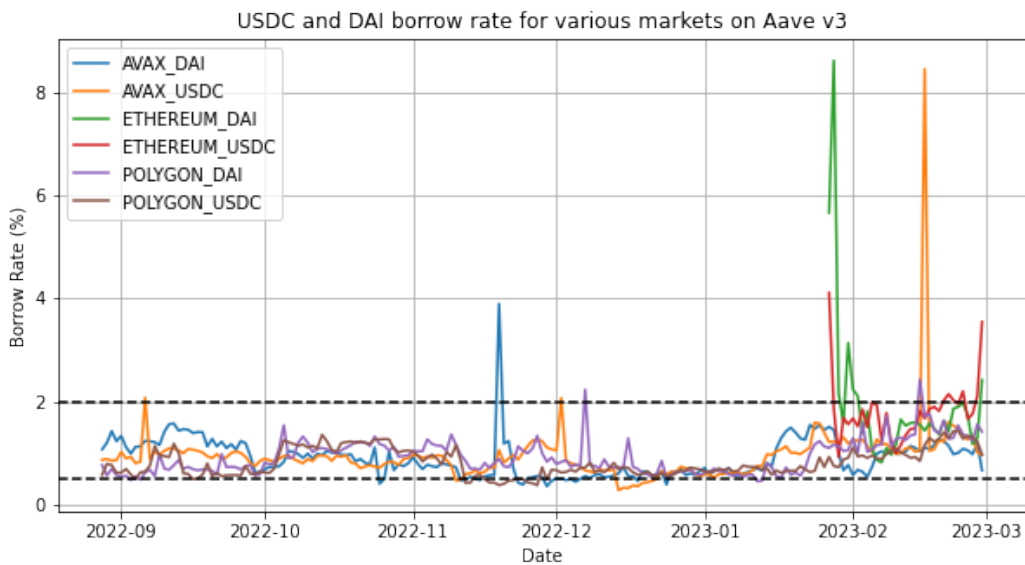


Figure 4: 6m historical average borrow rates for USDC and DAI across Avalanche, Ethereum, and Polygon. These were queried from the Aave API itself.

2.4 Fat Tail Probabilities

Accurately estimating probabilities of fat tail events (sometimes referred to as black or grey swans) from historical data is a markedly difficult task. Such estimates, particularly those using VaR or equivalent methods, are overwhelmingly prone to underestimating said probabilities [AT]. More modern attempts at predicting black swan events, such as extreme value theory, similarly underestimate the probability of fat tail events [Talb]. This leads to incorrectly justifying high debt ceilings. Frequently cited reasons for the difficulty in predicting black swans are lack of normality [Tala], lack of ergodicity [PG] [NW], and lack of historical data to train simulation models [Pau]. See Appendix B for more details on these concerns.

Instead of estimating the probability p , we propose a constraint for containing the losses of a catastrophic devaluation event inspired by standard banking regulation. This assumes the catastrophic devaluation event occurs with a probability of 1 and ensures a small percentage of the SM suffices to cushion the bad debt incurred. A good example of a potential fat tail event is the de-pegging of stablecoin Tether (USDT). Predicting the likelihood of such an event from historical data is unlikely to yield an accurate result. We argue we must ensure the solvency of the protocol were such an event to occur. Furthermore, we provide 3 scenarios of asset “riskiness” to understand the expected reserves from listing an isolated asset, assuming a conservative loss parameter L .

2.5 Minimum Risk Capital

In the realm of financial institutions, stress testing serves as a crucial tool to assess and ensure solvency even in exceptionally adverse conditions. To facilitate this process, the Federal Reserve, in accordance with the Dodd-Frank Act Stress Tests (DFAST) guidelines [Resa], furnishes financial institutions with a comprehensive set of macroeconomic scenarios (DFAST Scenarios, 2023). These scenarios provide a standardized framework for financial institutions to model various risk factors and evaluate their profit and loss (PnL) metrics under each designated scenario. Annually, financial institutions disclose the results of these stress tests, shedding light on their resilience and ability to withstand challenging market conditions.

non-stables have higher reserve factors) we arrive at approximately \$10M in reserve collections, close to the \$9.46M annualized revenue from token terminal [Ter]

Specifically, under the severely adverse scenario outlined by the Federal Reserve, bank-holding companies with significant trading operations are required to demonstrate sufficient capital holdings to absorb the default of their largest counterparty (DFAST Design, date). Notably, this default event is assigned a probability of 1, signifying a deterministic occurrence within the stress testing framework. By imposing this requirement, regulators aim to ensure that financial institutions possess adequate capital buffers to weather the potential default of their most prominent counterparty [Resc] under severe stress conditions. This default occurs with a probability of 1.

We can apply a similar principle when managing Aave’s exposure to isolated assets. Instead of estimating the probability of a catastrophic event, we instead manage the risk that it will occur with a probability of 1 upon listing the asset. We then ensure that the worst-case loss due to this event can be absorbed by some predetermined portion of the protocol’s safety module⁶. We propose the following constraint on setting the debt ceiling for risky assets:

Constraint 1: Ensuring Aave’s solvency during severely adverse market conditions.

Let $d_{a,m}$ denote the debt ceiling for asset a in market m , with a conservative loss parameter $L_{a,m}$. We are interested in managing the risk that bad debt approaches $\sum_{m \in M} d_{a,m} L_{a,m}$, denote this as Aave’s worst-case exposure to asset a across all markets M . Let Ω be the current size of the safety module. Assuming the debt ceiling and loss parameter are shared across markets for each asset (a benign assumption given cross-chain arbitrage), the following must be true:

$$0.01\Omega \geq \sum_{m \in M} d_a L_a \tag{5}$$

Notice that listing a as an isolated asset in more markets implies a greater systemic risk to the protocol.

Notice the protocol is listing multiple isolated assets. These assets likely exhibit material correlations and the mass liquidation of one asset might lead to the mass liquidation of other assets. Therefore, it is prudent to choose a small percentage of the safety module as a constraint but defer to the risk appetite of the community. Setting the limit to 1% for any isolated asset is a naive approach. Some risky assets exhibit greater [il]liquidity than others, so we might assign them a looser limit. As an initial risk-off approach, we find 1% to be a reasonable limit for all assets. As we continue to develop a more sophisticated risk management framework for isolated assets, we will expand upon this constraint.

2.6 Stress Testing our Debt Ceiling

Debt ceilings were previously set by considering an adversarial approach: i.e. how small does the debt ceiling need to be to minimize the profitability of an attack [Gau]. This differentiates two types of isolation mode assets: those whose prices are more easily manipulated (small-cap altcoins), and stablecoins. Attacks using stablecoins are more difficult since their prices are harder to manipulate (i.e. USDT) and usually have much larger market caps.

We set up our stress testing framework according to the average revenue parameter from section 2.3, and the conservative loss parameter in section 2.2. Recall that we are maximizing the expectation from equation 2. A brief sensitivity analysis of the optimal debt ceiling to the asset’s perceived riskiness and the stablecoin borrow rate is shown in figure 6.

Choosing a riskiness classification for an asset is analogous to determining the probability of a black swan. Simulations that rely on assumptions of asset prices and liquidity can be helpful, but are unlikely to accurately predict something like a USDT de-peg. That is, information on Tether’s market cap, supply, liquidity, etc., does not indicate anything meaningful about the quality of Tether’s reserves. Therefore, choosing riskiness, and therefore the debt ceiling, for an asset, will boil down to the community’s risk appetite and strategic goals moving forward. The stress test results above are meant to provide context on potential risks to guide the community’s decisions. Below we provide a possible rule of thumb for how the Aave community could think of this perceived risk:

⁶This is akin to the CET1 requirements (risk capital requirements) imposed by regulators on large financial institutions [Resb]. Regulators require financial institutions to hold a certain amount of high-quality assets (such as common stock) to absorb any losses from a percentage of their risk-weighted assets.

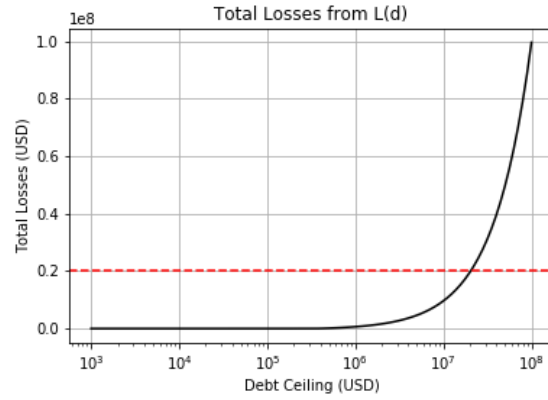


Figure 5: Gross losses given our loss curve $L(d)$. The red dashed line indicates the safety staking module constraint given a safety staking module worth $400M$ USDC.

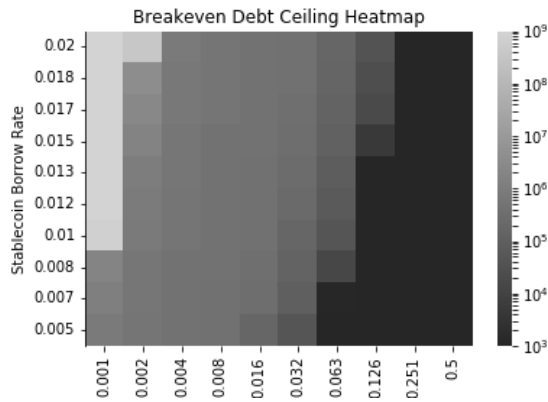


Figure 6: The recommended debt ceiling under various rate and risk scenarios (using the LUNA loss curve).

1. High risk assets: $p = 3\%$
2. Medium-high risk assets: $p = 1\%$
3. Medium risk assets: 0.2%

2.7 Long-Term Customer Acquisition and Retention

Although the non-negative profit constraint can help the protocol determine its debt ceilings in the long run, it may be less applicable to the growth-stage Aave protocol. Many growth-stage companies engage in product verticals that lose money in the short-term, yet have a lasting positive impact on userbase growth or alignment with other protocols. With this in mind, the community can reason about a reasonable expense that it would be willing to suffer in order to pursue these growth initiatives. It is beyond the scope of this paper to reason about a reasonable budget that the community should allocate toward such growth initiatives.

2.8 Recommendations

We provide example debt ceiling recommendations for the assets MKR, SNX, and YFI for Ethereum Mainnet. We emphasize that these debt ceilings only apply to Ethereum Mainnet, and these debt ceiling recommendations will be lower for other chains, given there is less liquidity (except for SNX since its liquidity is mostly on Optimism) and cross-chain arbitrage is not atomic. The liquidity ratio is as described in section 2.2.1, whereas the risk capital constraint is 1% of the Safety Module’s size as of February 2023. The breakeven amounts satisfy the zero-profit constraint from equation 2 at varying probabilities described in 2.6. These are meant to provide some context on how the perceived probability of a fat tail event affects the zero profit condition and highlight the minimal amount of reserves these assets generate relative to the risk they might pose during a shortfall.

	Liquidity Ratio	Risk Capital Constraint	High Risk Breakeven	Medium-High Risk Breakeven	Medium Risk Breakeven
UNI	9.6	7.5M	\$1.78M	\$3.91M	\$7.03M
MKR	3.4	5.2M	\$0.63M	\$1.38M	\$2.49M
SNX	0.6	4.2M	\$0.12M	\$0.26M	\$0.46M
YFI	0.2	4.1M	\$0.04M	\$0.09M	\$0.15M

Table 2: Results for select tokens on Ethereum mainnet.

We have sorted the tokens in order of on-chain liquidity on Ethereum mainnet, and Uniswap V3 pools. Notice that more liquid tokens generate higher liquidity ratios, and their loss parameters get adjusted accordingly. The perceived riskiness (i.e. the probability of a fat tail event) is much more complicated, and we provide some context on how the zero-profit debt ceiling scales with this fat tail risk. Notice that we have not included Tether in this table, since it does not pass our assumption in section 2.2.1 that its liquidity is distributed similarly to the LUNA liquidity during significant market stress. Furthermore, USDT is much more liquid than any of the tokens we have mentioned, with a lot of the liquidity concentrated on Curve.

2.9 Takeaways

1. We take a conservative approach to estimate the expected loss from an isolation mode asset given its debt ceiling.
2. We adjust this loss parameter according to the on-chain liquidity over relevant time periods, using TVL as a proxy for liquidity.
3. Similar to common risk and regulation practices in financial institutions, we provide a framework for setting a minimum risk capital requirement for isolated assets using Aave’s safety module. We propose

that the protocol’s worst-case exposure to any isolated asset across all markets be at most 1% of the safety module’s current size. This caps the potential debt ceiling gleaned from our stress test results.

4. Using this estimated loss, we then find a debt ceiling for which the expected revenue to the protocol’s reserves is equal to the expected loss from a catastrophic event. This breakeven point implies the losses from the catastrophic event can be paid for by the protocol’s reserves.
5. We provide a simple stress test of this breakeven point under different riskiness assumptions (source fig) to guide the community on setting the debt ceiling for 4 relevant assets: UNI, MKR, SNX, and YFI.

3 Discussion

We believe that this approach to setting the debt ceiling is reasonable, however it is likely not optimal. There are a few levers implicit in our methodology that could either relax or tighten the bounds we provide for listing risky assets. These levers, such as the loss parameter formulation, are discussed below.

3.1 Where our methodology excels.

This methodology for setting the debt ceiling is reasonable due to the fact that it covers the most obvious constraint: ensure that the protocol expects to not lose money from listing a risky asset in isolation mode. The validity of this constraint should be self-evident. However, requiring a non-negative profit expectation does not account for hard-to-quantify benefits of listing risky assets, as described in 2.7. Loosening this requirement is akin to the protocol expecting to pay for customer acquisition by listing more risky assets.

Our methodology is also quite simple, and thus hopefully straightforward for community members to understand. We have effectively reduced the debt ceiling problem into two subproblems: an estimation of the asset’s catastrophe risk, and an estimation of the reserves that an asset will generate for the protocol. Over time, we can refine models to estimate these two quantities with higher precision.

Furthermore, our methodology is very conservative. As an initial framework for understanding isolated asset risk, we aimed at providing a very risk-off approach to listing isolated assets or raising their debt ceilings. This can be both a pro, in terms of ensuring the protocol’s solvency, and a con, in terms of quickly growing the protocol. In section 2.2 we assume the entire debt ceiling is in default, and whatever isn’t liquidated in 1 transaction is a protocol loss. Our approach is also based on WETH/LUNC liquidity, which might be representative of less liquid risky assets, such as SNX, but is likely not representative of how a USDT liquidation would play out. We could imagine improving this framework by simulating these liquidations using agent based simulations, or by considering multiple liquidating transactions over multiple blocks. Both of these methods would provide a more accurate loss parameter as a function of the debt ceiling, which might justify relaxing the debt ceiling constraints.

3.2 How our methodology can be improved in the future.

3.2.1 Zero-Profit

As mentioned, our framework for setting the debt ceiling entails a zero-profit point for the protocol. In the future, it would be better for the protocol to perform a numerical optimization over the space of possible debt ceilings. For instance, if we let d_{max} be the zero-profit debt ceiling, in the future we might do the following optimization:

$$\begin{aligned} \max_d \quad & U(d) \\ \text{s.t.} \quad & d \leq d_{max}, \end{aligned} \tag{6}$$

where $U : \mathbb{R}^+ \rightarrow \mathbb{R}$ is a utility function defined over the space of possible debt ceiling parameters. If we wanted to maximize profit, we could set the utility functions to

$$U(d) = \mathbb{E}[\text{profit} \mid \text{debt ceiling} = d].$$

If we wanted to maximize a risk-adjusted profit metric, we could utilize the following utility function:

$$U(d) = \frac{\mathbb{E}[\text{profit} \mid \text{debt ceiling} = d]}{\text{Std}[\text{profit} \mid \text{debt ceiling} = d]}.$$

We leave utility function optimizations to future work.

3.2.2 Safety Module Constraint

Requiring that the sum of an isolated asset's debt ceiling across all markets be less than 1% of the safety module is a reasonable first pass. However, depending on the risk appetite of the community this constraint can be tightened or loosened. There are two key areas of research that could help ensure Aave always holds sufficient risk capital: (1) factoring the variance of the safety module's size into the debt ceiling limits, and (2) modifying the 1% constraint as more risky assets are listed and their contagion effects (eg. price correlation) are better understood. For example, depending on the safety module's value, we may be more or less willing to set a high debt ceiling. A scenario-based analysis, where we consider various trajectories of the value of the safety module, would allow us to determine debt ceilings in such a way that would lead to us lower probability of losses that exceed a fixed percentage of the safety module's value.

A Uniswap v3 Mechanics

A Uni v3 pool allows liquidity providers to provide liquidity between specified ticks (prices) in the v3 pool. Denote the space between two consecutive ticks as a bucket. The bucket has a specified price, the left tick, and a certain amount of liquidity contained within it [Ada], illustrated in Figure 7.

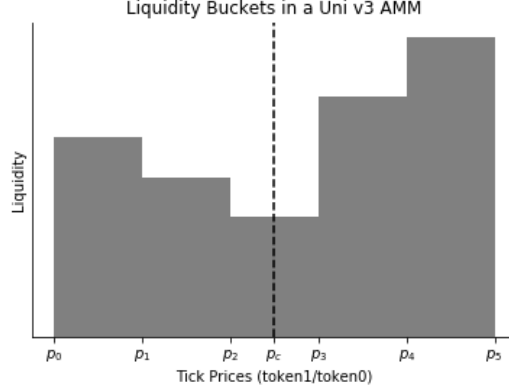


Figure 7: The liquidity buckets for an example uni v3 pool. Each p_i indicates the tick price of the bucket to its right. The tick p_c represents the current pool price.

The liquidator’s transaction begins at the current pool price p_c . Every marginal unit of LUNA swapped into the pool pushes the current price p_c to the right on the graph in exchange for some ETH. Let δ_y denote the amount of LUNA swapped into the bucket, and δ_x be the corresponding ETH received. Let p_i denote either p_c (for the first bucket) or the left tick of a given bucket, then:

$$\delta_x = \lambda_i \left(\frac{1}{\sqrt{p_i} + \frac{\delta_y}{L}} - \frac{1}{\sqrt{p_i}} \right) \quad (7)$$

Where λ_i is the bucket’s liquidity, [Ada]. In each bucket, the liquidator can swap up to $\delta_y \leq L(\sqrt{p_{i+1}} - \sqrt{p_i})$ LUNA before pushing the price over to the next bucket [Ada]. At every bucket, the liquidator calculates her marginal profit from swapping up to the bucket’s limit in LUNA. For ease of exposition we denominate costs, revenues, and profits in units of LUNA. Her cost is the marginal LUNA swapped into the bucket, δ_y , whereas her revenue is the LUNA value that she receives from the protocol upon repaying the USDC debt. She receives $(1 - \gamma_{WETH,LUNA})\delta_x$ in ETH from the LUNA swap (where $\gamma_{WETH,LUNA}$ is the WETH/LUNA swap fee, = 0.01) and swaps it into $(1 - \gamma_{WETH,LUNA})(1 - \gamma_{WETH,USDC})p_{USDC,ETH}\delta_x$ USDC (where γ is the USDC/ETH pool fee, = 0.0005). She repays this much USDC to the protocol and receives $(1 + LB)(1 - \gamma_{WETH,LUNA})(1 - \gamma_{WETH,USDC})p_{USDC,ETH}p_{LUNA,USDC}\delta_x$ in LUNA, where LB is the collateral’s liquidation bonus⁷. Using equation 7, we can express the liquidator’s profits in a given bucket in terms of the LUNA swapped, δ_y . For simplicity, let $\alpha = (1 - \gamma_{WETH,LUNA})(1 - \gamma_{WETH,USDC})p_{USDC,ETH}p_{LUNA,USDC}$, a constant, then:

$$\text{Profit}(\delta_y) := (1 + LB)\alpha\lambda_i \left(\frac{1}{\sqrt{p_i} + \frac{\delta_y}{L}} - \frac{1}{\sqrt{p_i}} \right) \quad (8)$$

As shown in Figure 8, the liquidator’s profit curve for a given bucket is concave. We find the profit maximizing LUNA for each bucket by maximizing equation 8:

$$\delta'_y = \min \left(\lambda_i \cdot \left(\sqrt{\alpha \cdot (LB + 1)} - \sqrt{p_i} \right), \lambda_i (\sqrt{p_{i+1}} - \sqrt{p_i}) \right) \quad (9)$$

⁷We assume a bonus of 15%.

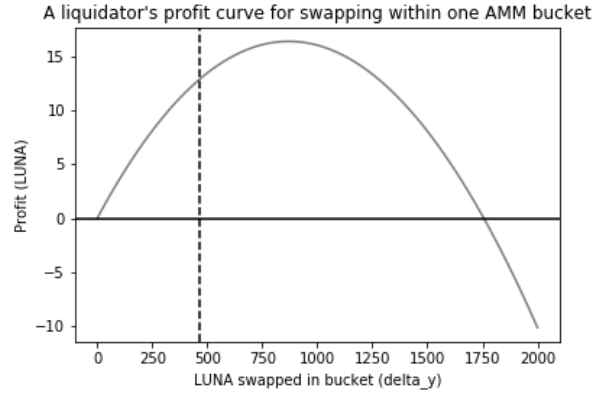


Figure 8: A sample profit curve for a uni v3 bucket in block 14740860. The dashed line denotes the maximum LUNA that can be swapped before pushing the price to the next bucket.

The liquidator will continue to swap LUNA for ETH as long as $\delta'_y > 0$. Profits monotonically increase the more LUNA is swapped, so once $\delta'_y \leq 0$ for bucket $i = j$, swapping LUNA in any bucket $i > j$ would not be profitable. Furthermore, the liquidator would only continue to liquidate so long as there is USDC debt to be repaid, and LUNA collateral to be claim. In usual circumstances, the liquidator could only claim up to half the defaulting collateral in a single liquidation. However, if the health factor dips below 0.95, the entire debt ceiling can be repaid and the entire defaulting collateral may be claimed [Lab]. Since we are taking a conservative approach, we assume a worst case scenario where health dips to 0.95.

B Fat Tail Probabilities

There are a number of concerns regarding such black swan events, which lead statistical models to often underestimate the impact, and the probability, of black swan events. We explain some of them below:

1. **Normality:** Probability distributions in financial systems may behave normally around the mean, but rarely do so at the tails. When managing risk for black swan events, using models that presume a Gaussian distribution will underestimate the likelihood of such events. To understand black swans, we need fatter tailed distributions, such as a Mandelbrotian distribution. See [Tala] for a concise example for how to reject the assumption that a distribution is Gaussian once a black swan has been observed.
2. **Ergodicity:** Using expectations could be misleading as the cost/benefit is materially asymmetric at the margins. This asymmetry arises from the existence of “absorbing states”, and is often referred to as the “risk of ruin”. An example of an absorbing state in a financial system is bankruptcy: an agent that reaches bankruptcy can no longer play the game. The existence of absorbing states leads to irrational cost benefit analysis results. In the context of Aave, insolvency could be understood as an absorbing state: if bad debt exceeds the protocol’s assets, users lose their ability to withdraw and, in turn, the protocol loses its users (potentially permanently). Most financial models, such as modern portfolio theory, assume the lack of such absorbing states (ie. they assume ergodicity), which breaks apart at the margins. To address this concern, the absorbing state must be avoided (ie. we say no to potential fat tail risks). In this case, we can address this with hard constraints on the debt ceiling, for example, by requiring the sum of debt ceilings of correlated assets to never exceed a certain percentage of the safety module (maybe 100%).
3. **Data:** There is an insufficient lookback period for a historical analysis, especially in an emerging market such as crypto. Extrapolating statistical variables on nascent tokens is bound to suffer from large statistical uncertainty.

These problems naturally arise when trying to manage the risk of highly improbable events, and generally lead to the underestimation of their probabilities. Instead of trying to estimate the probability of the black swan event happening, we will adopt an approach similar to that of regulated financial institutions, who face similar concerns.

C Authors

Omer Goldberg `omer@chaoslabs.xyz`

Yonatan Haimowitz `haimo.yonatan@chaoslabs.xyz`

References

- [Ada] Hayden et al Adams. “Uniswap v3 Core”. In: (). URL: <https://uniswap.org/whitepaper-v3.pdf>. accessed: 02/27/2023.
- [AT] Michael Adams and Barry Thornton. “Black Swans and VaR”. In: (). URL: <https://www.aabri.com/manuscripts/131653.pdf>. accessed: 02/23/2023.
- [Coma] Aave Companies. *Avalanche*. URL: <https://docs.aave.com/risk/asset-risk/avalanche>. accessed: 02/24/2023.
- [Comb] Aave Companies. *Risk Methodology*. URL: <https://docs.aave.com/risk/asset-risk/methodology>. accessed: 02/24/2023.
- [FH22] Emilio Frangella and Lasse Herskind. “Aave V3 Technical Paper”. In: (2022).
- [Gau] Gauntlet. *ARC Aave v3 Ethereum Deployment Assets and Configurations*. URL: <https://governance.aave.com/t/arc-aave-v3-ethereum-deployment-assets-and-configurations/10238/22>. accessed: 02/27/2023.
- [HG23] Yonatan Haimovich and Omer Goldberg. “Aave V3 Risk Parameter Methodology”. In: (2023).
- [Lab] Uniswap Labs. *Uniswap Liquidation Logic*. URL: <https://github.com/aave/aave-v3-core/blob/master/contracts/protocol/libraries/logic/LiquidationLogic.sol#L49-L68>. accessed: 02/27/2023.
- [NW] Duncan Needham and Julius Weitzdorfer. *Extremes*. accessed: 02/23/2023.
- [Pau] Pauljlei. *[ARC] Aave V3 Ethereum Deployment: Assets and Configurations*. URL: <https://governance.aave.com/t/arc-aave-v3-ethereum-deployment-assets-and-configurations/10238/22>. accessed: 02/24/2023.
- [PG] Ole Peters and Murray Gell-Man. *Evaluating gambles using dynamics*. URL: <https://arxiv.org/pdf/1405.0585.pdf>. accessed: 02/23/2023.
- [Resa] Federal Reserve. *2023 Stress Testing Scenarios*. URL: <https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20230209a1.pdf>. accessed: 02/24/2023.
- [Resb] Federal Reserve. *Large Bank Capital Requirements*. URL: <https://www.federalreserve.gov/publications/large-bank-capital-requirements-20210805.htm>. accessed: 02/24/2023.
- [Resc] Federal Reserve. *Policy Statement on the Scenario Design Framework for Stress Testing*. URL: <https://www.govinfo.gov/content/pkg/FR-2017-12-15/pdf/2017-26858.pdf>. accessed: 02/24/2023.
- [Tala] Nassim Nicholas Taleb. *Book screenshot on rejecting normality on fat tail events*. URL: <https://twitter.com/nntaleb/status/1619714529172291585/photo/1>. accessed: 02/23/2023.
- [Talb] Nassim Nicholas Taleb. *Extreme Value Theory: Fuhgetaboutdit*. URL: <https://www.fooledbyrandomness.com/EVT.pdf>. accessed: 02/23/2023.
- [Ter] Token Terminal. *Aave (AAVE)*. URL: <https://tokenterminal.com/terminal/projects/aave>. accessed: 02/01/2023.