

Stress Testing Aave’s Supply Cap Parameter

Chaos Labs

March 2023

Abstract

We propose a stress-testing framework for setting the Aave supply cap parameter for all Aave v3 tokens. We search for supply caps that keep protocol losses at sustainable levels under stressed agent-based simulations. We stress test the accrual of bad debt on Aave against hypothetical (but tangible) scenarios. To construct our stress scenarios and their corresponding price and liquidity trajectories, we glean data from recent stablecoin depegs and token devaluations. We then perform a sensitivity analysis on our results, observing how protocol losses change if we exaggerate the rate at which price and liquidity deteriorate. Our stress testing framework can be extended to other risk parameters, such as debt ceilings and borrow caps.

1 Introduction

Aave v3 introduces various risk management tools and parameters. These parameters are currently set by simulating stressed Value at Risk (VaR) using an agent-based simulation approach. These stressed VaR methodologies extrapolate risk from aggregating historical data based on statistical indicators such as volatility and drift. Alternatively, stress tests project an institution’s risk during severely adverse market conditions, often muted or pruned from most stressed VaR methods [4]. Stress tests are particularly useful for financial institutions to ensure they survive certain idiosyncratic events [1]. By running well-designed stress tests, we can gauge a protocol’s potential exposure to various hypothetical (but tangible) scenarios, for example:

1. A sudden 99% stablecoin depeg.
2. A sudden 25% ETH devaluation.

These scenarios, while seemingly unlikely, have been observed a handful of times in recent history. This paper will glean price and liquidity trajectories from historical examples of various stress scenarios. We will then map Aave’s listed tokens to these price and liquidity trajectories, and, using Chaos Labs’ agent-based simulation framework [2], we obtain the expected losses to the protocol. Ultimately, the goal of this process is to ensure Aave’s solvency in spite of material market turmoil.

We will use this stress test approach to set supply caps for Aave’s listed tokens. This is a reasonable framework because these rare events happen significantly more frequently than expected if their probability distributions were normal. Supply caps are one crucial parameter in mitigating this fat tail risk. The stress testing framework presented in this paper may also be applied more generally to stress test various other elements of the Aave protocol, including the isolated asset debt ceiling [3].

1.1 Liquidity Risk

We will use stress tests to measure Aave’s liquidity risk with respect to each of its listed tokens. We provide an Aave-specific definition of liquidity risk below:

Definition 1 (Liquidity Risk). *The risk is that a token’s collateral supply cannot be sold profitably for the outstanding debt token once a position becomes under-collateralized. This leaves liquidators without incentive to liquidate underwater accounts, allowing the gap between collateral and debt value to grow. This occurs particularly due to exchange slippage. Aave is left with bad debt either if prices never converge (missed liquidations) or if liquidations begin to occur when account margins are significantly below their liquidation threshold (adverse liquidations). Adverse liquidations lead to the bad debt since the entire collateral may be profitably liquidated by a liquidator without paying back the outstanding debt in full.*

Liquidity risk depends on the ratio between how much collateral must be liquidated and how much liquidity is available on decentralized exchanges (DEXes)¹. The supply cap, as we discuss in Section 3, allows us to cap how much of any token can be supplied on Aave and, therefore, how much can be used as collateral. Therefore, stress testing the supply cap is a liquidity risk test, informing us of how much collateral can be supplied on Aave before it becomes too costly to liquidate under duress.

1.2 Reference Events

For each stress scenario, we will consider one relevant *reference event* from which we base our simulations. We define these reference events below and consider including more reference events and constructing synthetic reference events in Sections 4.2 and 4.3.

Definition 2 (Reference Event). *A reference event is a historical example of its stress scenario. These reference events will be used to obtain price and liquidity trajectories. By pulling decentralized exchange data, we can construct the rate of returns on prices and liquidity and map those onto each token’s current on-chain price and liquidity. We will limit ourselves to reference events for which we can obtain on-chain data to simulate potential liquidations.*

2 Stress Testing

Stress testing all tokens against the same stress scenarios might lead to unreasonable recommendations. For example, stress testing BTC against a LUNA-like crash would lead to extremely conservative supply caps on WBTC, using an extremely unlikely scenario (BTC goes to zero in a few days). At the logical limit, we wouldn’t stress test any token against the scenario that it goes to 0 instantaneously. Finding the perfect stress test is a complicated task. We aim to consider concrete examples informing us of Aave’s various risks, so we can appropriately mitigate them. From the IMF’s guidelines on designing stress tests:

In discussions on designing stress tests, too much is often made of establishing the “right” scenario. [...] In practice, assessing such consistency is tricky because the scenarios should also be exceptional (but plausible). [...] One approach is to choose a concrete, extreme historical scenario (e.g., the East Asian

¹Liquidators interact primarily with DEXes, so liquidity concentrated on centralized exchanges does not meaningfully reduce Aave’s liquidity risk

crisis of 1997) and calculate the impact of repeating such a scenario (or an adaptation of such a scenario) in the present situation of the banking system. The main advantage is that historical scenarios are easy to communicate and implement. Also, they are plausible because such a situation actually happened. Their main disadvantage is that past crises may not be good models for future crises [1].

In designing our stress tests, we consider stablecoins and non-stablecoins separately. We create a baseline stress test for each category, described in Sections 2.1.2 and 2.1.4, respectively. For riskier assets, we construct more severe stress tests: for stablecoins, we simulate a UST-esque collapse, whereas, for non-stablecoins, we consider UST’s sister token, LUNA. These four stress scenarios provide an initial comprehensive suite to stress test Aave’s supply cap for each token. It could be extended to other risk parameters, such as the isolated asset debt ceiling.

2.1 The Stress Scenarios

Given computational limitations, we limit the time window of each reference event to a couple of weeks. We query Dune Analytics DEX data² to obtain the primary sources of DEX liquidity for each reference token. We cannot explain all the liquidity (e.g., OTC trades, various smaller DEX pools), making our simulations slightly more conservative. However, since liquidators likely do not have immediate access to OTC markets for liquidations and are not integrated with smaller and less liquid DEXes, we are likely getting a reasonable estimate. We consider additional stress tests for specific risks, such as a CRV-esque short-seller attack or an MKR-esque governance token auction, in Section 4.2.

2.1.1 Severe Stablecoin Depeg

The severe stablecoin depeg scenario tests for a complete 99% depeg of a stablecoin. It applies primarily to small-cap or algorithmic stablecoins, which intuitively are at the highest risk of such a dramatic devaluation. This also serves as an initial conservative test for newly listed stablecoins or those the community perceives have out-sized risk. For example, we might want to test FRAX or LUSD against this scenario. We test FRAX because it is partially algorithmic and currently undercollateralized. We test LUSD because it is a relatively young stablecoin with a smaller market cap.

Reference Event: The May 2022 UST collapse. The collapse of Terra’s USD stablecoin was crypto’s largest stablecoin collapse. It provides a useful worst-case scenario for any of the riskier stablecoins on Aave. Particularly, we simulate protocol losses if any partially or fully algorithmic stablecoin, such as FRAX, experiences a similar collapse.

- Start Time: May 7, 2022 12:00:00 AM³
- End Time: May 15, 2022, 12:00:00 AM

2.1.2 Moderate Stablecoin Depeg

The moderate stablecoin depeg tests for an approximately 20% depeg of a large stablecoin. As discussed in later sections, this drawdown can be scaled to consider a more severe depeg. The primary risks we consider in this scenario are, for example, a coin issuer’s temporary loss of access to reserve assets (such as the SVB/USDC event in March 2023) and potential duration/interest rate risk in reserve assets. Either risk

²<https://dune.com/queries/2249304/3686312>

³We begin all our simulations at 12AM GMT. This does not have a material impact on our simulations as long as we select an appropriate start day.

Project	Version	Contract	USD Vol	Pct Vol
Curve	Factory V1 Meta	0xceaf7747579696a2f0bb206a14210e3c9e6fb269	\$5,503,659,362	80.12%
Uniswap	3	0x7a5ae802895d5f90b6edbafc870fd348fba2a3d2	\$438,536,848	6.38%
Uniswap	3	0xcc527260a3f3adf15698689da926aa72a7f1409b	\$326,361,323	4.75%
Sushiswap	1	0x9a0cc6791a5409ce3547f1f1d00e058c79d0a72c	\$287,170,504	4.18%
Uniswap	3	0xa87b2ff0759f5b82c7ec86444a70f25c6bffcbbf	\$167,331,507	2.44%
Uniswap	3	0x12d6867fa648d269835cf69b49f125147754b54d	\$92,012,097	1.34%

Table 1: UST DEX Liquidity during the UST reference event.

could lead to a “bank run” on the issuer and might lead the token to trade at a discount on DEXes for some period of time.

Reference Event: The USDC depeg during the bank run on Silicon Valley Bank in March of 2023. This was one of the largest depeg observed by one of the largest circulating stablecoins. It occurred due to Circle’s temporary loss of access to reserve assets, which was resolved by the FDIC’s commitment to ensure all SVB depositors. We consider how exaggerated versions of this event could impact Aave’s potential losses.

- Start Time: March 10, 2023, 12:00:00 AM
- End Time: March 14, 2023, 12:00:00 AM

Project	Version	Contract	USD Vol	Pct Vol
Curve	Regular	0xbebc44782c7db0a1a60cb6fe97d0b483032ff1c7	\$6,856,778,994	25.5%
Uniswap	3	0x88e6a0c2ddd26feeb64f039a2c41296fcb3f5640	\$6,441,732,449	23.96%
Uniswap	3	0x7858e59e0c01ea06df3af3d20ac7b0003275d4bf	\$4,661,573,230	17.34%
Uniswap	3	0x3416cf6c708da44db2624d63ea0aaef7113527c6	\$1,847,555,181	6.87%
1inch Limit Order Protocol	3	0x1111111254eeb25477b68fb85ed929f73a960582	\$1,294,104,939	4.81%
Curve	Regular	0xa5407eae9ba41422680e2e00537571bcc53efbfd	\$1,269,449,147	4.72%
Uniswap	3	0xe0554a476a092703abdb3ef35c80e0d76d32939f	\$1,103,516,084	4.1%
Uniswap	3	0x5777d92f208679db4b9778590fa3cab3ac9e2168	\$1,002,694,011	3.73%
Uniswap	3	0x6c6bc977e13df9b0de53b251522280bb72383700	\$606,836,794	2.26%
Uniswap	3	0x8ad599c3a0ff1de082011efddc58f1908eb6e6d8	\$380,745,831	1.42%
DODO	1	0xc9f93163c99695c6526b799ebca2207fd7d61ad	\$325,600,035	1.21%

Table 2: USDC DEX Liquidity during the USDC depeg event.

2.1.3 Severe Token Drawdown

The severe token drawdown scenario is a worst-case stress test of non-stablecoins. We simulate an approximately 99% devaluation of the token from the beginning of the reference event to the end of the reference event. This is applicable to tokens perceived by risk providers or the Aave community as extremely risky. Risks include extreme price manipulation (such as occurred with the MNGO token exploit), a smart contract exploit, or a LUNA-esque death spiral.

Reference Event: The collapse of the LUNA token.

- Start Time: May 7, 2022, 12:00:00 AM
- End Time: May 15, 2022, 12:00:00 AM

Project	Version	Contract	USD Vol	Pct Vol
Sushiswap	1	0x88172e5d79fe75c7aed1453e89ff5d741cfa4ca7	\$371,621,067	78.35%
Uniswap	2	0x60a39010e4892b862d1bb6bdde908215ac5af6f3	\$39,882,859	8.41%
Uniswap	3	0x16b70f44719b227278a2dc1122e8106cc929ecd1	\$31,142,367	6.57%
Uniswap	3	0xf93d138e087b454b4051496346bcc61e4f8f99d4	\$23,503,524	4.96%

Table 3: LUNA DEX Liquidity during the LUNA reference event.

2.1.4 Baseline Token Drawdown

This is the primary stress test for most tokens. We test tokens against the drawdown in prices and on-chain liquidity during the collapse of the Terra ecosystem in May of 2022. Terra’s collapse caused a roughly 35% drawdown in both ETH and BTC throughout May 2022. We look at ETH price and liquidity trajectories throughout this reference event. This provides us with a benchmark for the “shape” of price and liquidity trajectories under significant market stress. We then map it onto each of the tokens we are stress testing to understand how depreciation in that token’s price and liquidity might affect the Aave protocol. Since each token has a unique risk profile, we consider scaling the rate at which price and liquidity deteriorate by different factors. This scaling is often referred to as sensitivity analysis in risk management and allows us to differentiate between adverse market conditions and severely adverse market conditions. It is discussed in Section 4.1.

Reference Event: The Terra collapse.

- Start Time: May 5, 2022, 12:00:00 AM
- End Time: May 13, 2022, 12:00:00 AM

Project	Version	Contract	USD Vol	Pct Vol
Uniswap	3	0x88e6a0c2ddd26feeb64f039a2c41296fcb3f5640	\$9,906,784,399	36.58%
Uniswap	3	0x4585fe77225b41b697c938b018e2ac67ac5a20c0	\$2,396,819,226	8.85%
Uniswap	3	0x8ad599c3a0ff1de082011efddc58f1908eb6e6d8	\$1,812,266,353	6.69%
Curve	Regular	0xd51a44d3fae010294c616388b506acda1bfaae46	\$1,740,517,132	6.43%
Uniswap	3	0x11b815efb8f581194ae79006d24e0d814b7697f6	\$957,161,685	3.53%
Uniswap	3	0x4e68ccd3e89f51c3074ca5072bbac773960dfa36	\$954,101,625	3.52%
Uniswap	2	0xb4e16d0168e52d35cacad2c6185b44281ec28c9dc	\$620,965,404	2.29%
Uniswap	3	0xcbcdf9626bc03e24f779434178a73a0b4bad62ed	\$597,426,406	2.21%
Uniswap	3	0xac4b3dacb91461209ae9d41ec517c2b9cb1b7daf	\$434,737,119	1.61%
Uniswap	3	0x7a5ae802895d5f90b6edbafcb870fd348fba2a3d2	\$423,467,112	1.56%
Sushiswap	1	0x397ff1542f962076d0bfe58ea045ffa2d347aca0	\$414,048,855	1.53%
0x Native	4	0xdef1c0ded9bec7f1a1670819833240f027b25eff	\$412,994,279	1.52%
Uniswap	2	0x0d4a11d5eeaac28ec3f61d100daf4d40471f1852	\$354,762,968	1.31%
Balancer	2	0x96646936b91d6b9d7d0c47c496afb3d6ec7b6f8	\$326,214,683	1.2%
0x Native	2.1	0x080bf510fcbf18b91105470639e9561022937712	\$299,528,562	1.11%

Table 4: ETH DEX Liquidity during the Terra collapse reference event.

2.2 Matching Tokens to Stress Scenarios

For example, we bucket each Aave Ethereum v3 token according to its most appropriate category in Table 5. In general, all tokens are considered for either the moderate stablecoin depeg or the moderate token drawdown. We use the scale factor to perform our sensitivity analysis: a higher scale factor means that the

simulated price and liquidity trajectories deteriorate at a greater speed than the referenced event (i.e., ETH and USDC). The scale factors in the table are for reference, ultimately, we will consider scale factors of (1%, 5%, and 10%) for all tokens and let the Aave community decide its risk appetite for each token. We label these scale factors as baseline, adverse, and severely adverse. This is further explained in Section 4.1.

We aim to stress test each token against a reasonable stress scenario. If we are too conservative, e.g. we stress test USDC against a 99% depeg, we would arrive at unreasonably low supply caps, damaging Aave’s ability to provide a useful service to users; we cannot fully eliminate all risk. The severe token drawdown and severe stablecoin depeg scenarios are useful to stress test new or particularly risky tokens, which we don’t believe apply to many assets currently listed on Aave’s Ethereum v3 market (but might apply to tokens on other markets, e.g. AaveGotchi on Polygon).

Token	Stress Scenario	Scale Factor
DAI	Moderate Stablecoin Depeg	5%
LUSD	Severe Stablecoin Depegs	10%
USDC	Moderate Stablecoin Depeg	5%
USDT	Moderate Stablecoin Depeg	5%
AAVE	Moderate Token Drawdown	10%
cbETH	Moderate Token Drawdown	10%
CRV	Moderate Token Drawdown	10%
LINK	Moderate Token Drawdown	5%
rETH	Moderate Token Drawdown	10%
WBTC	Moderate Token Drawdown	1%
WETH	Moderate Token Drawdown	1%
wstETH	Moderate Token Drawdown	10%

Table 5: Aave v3 Ethereum Tokens

2.3 Stress Scenario Summary

Below we describe some of our stress-testing approach’s assumptions, strengths, and weaknesses. They are discussed in further depth in Section 4.

1. **Stablecoins and Non-Stablecoins:** We assert that stablecoins pose unique risks to the protocol that must be tested separately from other tokens. We aim to capture the unique risk profiles of stablecoins relative to other types of tokens.
2. **Historical Predictive Power:** We stress test protocol losses according to prices, liquidities, and covariances observed during our reference events. However, historical events are likely not perfectly representative of future risks. This limitation is present in most standard risk management practices [1] [4]. To better understand the consequences of this assumption, we propose a sensitivity analysis framework in Section 4.1.
3. **Stress Scenario Selection:** Crypto is an emerging market; price and liquidity profiles change significantly over the years. We bias our selection of stress tests for more recent events. For example, we don’t test BTC against the Mt. Gox exploit, as this is unlikely to yield reasonable results relative to today’s BTC demand and liquidity. For this reason, we focus on events that occurred throughout 2022 and 2023.
4. **Ethereum Liquidity:** We observe the price and liquidity trajectories of each reference event by looking exclusively at Ethereum DEXes. These are primarily Uniswap v2 and v3 for most assets, and

Curve v1 and v2 for for stablecoins. L2 chains, such as Avalanche or Arbitrum, might experience worse liquidity deterioration as liquidators can't atomically interact with Ethereum liquidity venues (i.e. in one transaction), and there is less liquidity on L2-native DEXes.

5. **Additional Scenarios:** We consider additional stress scenarios in Section 4.2, such as an ETH liquid staking derivative stress scenario. We consider how to develop synthetic reference events for these scenarios, especially for those with insufficient historical reference events, in Section 4.3.
6. **Terra Luna:** Our stress tests significantly rely on data from the Terra ecosystem collapse. While we consider the price and liquidity trajectories of different tokens during this period, we might be overfitting the specific market behavior of this historical event. However, the Terra collapse offers many recent and extreme movements in price and liquidity that are useful in our stress tests. This would contrast to older or more mild stress tests, such as MakerDAO's auction of MKR on "Black Thursday" in 2020, or the recent short-seller attack on CRV by Avraham Eisenberg.

3 Setting the Supply Cap

We aim to demonstrate a methodology for setting the supply cap for a token on the Aave v3 protocol. Implicitly, we set the supply cap by considering how much of that supply is being used as collateral. We base our methodology on the premise that the supply cap's role in the Aave risk parameter suite is to limit the protocol's exposure to liquidity risk, particularly the slippage experienced by liquidators. If this slippage exceeds the liquidation bonus for liquidators, debts are not repaid. This liquidity risk can be mitigated by reducing Aave's collateral exposure to certain tokens. We determine the supply cap by measuring how increases in collateral exposure increase liquidity risk and therefore protocol losses.

3.1 Methodology

We use Chaos Labs' simulation engine to binary search for the supply cap. We state it in words and pseudocode in Algorithm 1. We apply this binary search for a supply cap under stressed price and liquidity trajectories, defined in the previous Section.

3.1.1 Supply Cap Search

In our search for an optimal supply cap, we assume users are elastic to the supply cap. When the supply cap for a token increases, say it doubles, users double their supply of that token and the amount borrowed against it. We instantiate vectors of user borrows and supplies, scale their supply and borrow of a particular token, and then play out our stress test. This effectively allows us to search for the optimal amount of collateral for the token in question and set its supply cap. We discuss this elasticity assumption, and other assumptions, in Section 3.1.2

Step 0 Initialize constants. For ease of exposition, we refer to the pool's underlying token by the name 'A', this is the token being supplied. Let:

- $S_{min} = 0$ be the minimum feasible supply cap.
- $S_{max} =$ be the maximum feasible supply cap. **We set this as the circulating supply of token A⁴.**

⁴This means we won't set supply caps greater than the current supply of the token, even if it means the expected loss is below our target loss K .

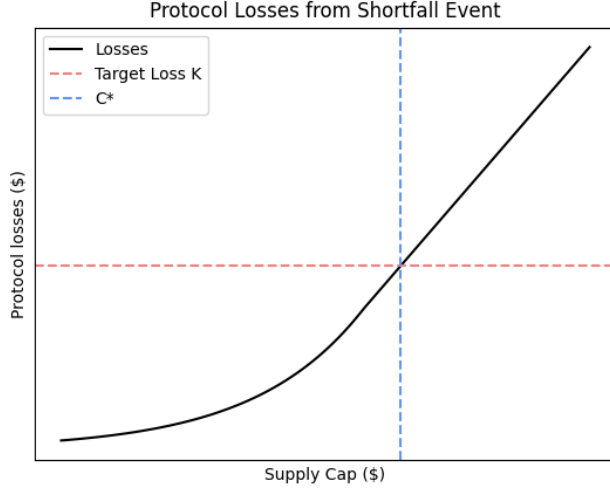


Figure 1: The protocol losses with respect to the supply cap is an increasing function that we do not know analytically. We simulate various supply caps to compute these losses and binary search between C_{min} and C_{max} to find the optimal supply cap, C^* , that leads to a protocol loss of K . Intuitively, we know protocol losses asymptotically approach a linear increasing function, $y = x$, once all liquidity is consumed and any additional collateral pledged leads to a complete loss of the outstanding token.

- $S_{cur} = \frac{S_{min} + S_{max}}{2}$, the initial cap we consider in our binary search.
- $K \geq 0$ be the target protocol losses in the case of a stress scenario, described in Section 3.2.
- ε be a ‘precision’ variable that determines the maximum distance from K we will accept for stopping the supply cap search⁵. **To ensure the program terminates, we also terminate if $S_{max} - S_{cur} < \varepsilon$.**

Step 1 Determine stressed price and liquidity trajectories:

1. Obtain the USD price and liquidity trajectories of the reference token during its reference event, defined in Section 2. We obtain these price and liquidity trajectories by querying on-chain data for relevant DEX liquidity venues (e.g., Section 2.1.2). For the reference token R , the price at time t during the stress event is denoted as $p_{R,t}$, and the liquidity as $L_{R,t}$.
2. Query the current (or recent average) price of our target token A , denote it as $p_{A,0}$.
3. Generate the stressed price trajectory of A as:

$$p_{A,t} = p_{A,0} \frac{p_{R,t}}{p_{R,0}}. \quad (1)$$

4. Similarly, we generate the stressed liquidity trajectory of A as:

$$L_{A,t} = L_{A,0} \cdot \frac{L_{R,t}}{L_{R,0}}. \quad (2)$$

⁵For instance, if $K = \$10,000$ and $\varepsilon = \$1,000$, then we are targeting a supply cap that will experience a loss of \$10,000 in a stress scenario, and we will stop searching for the supply cap once our on-chain simulation losses are within the range $[\$10,000 - \$1,000, \$10,000 + \$1,000] = [\$9,000, \$11,000]$.

Step 2 Agent-based simulations rely on the distribution of borrows and supply. We query the current distribution of borrows and supplies to the pool: \vec{s} and \vec{b} . Each is a two-dimension vector: $\vec{s} = \{\vec{s}_i\} = \{s_{i,j}\}$ where i denotes an address, \vec{s}_i denotes the tokens supplied by address i , and j denotes the token - e.g. $s_{0xabc,USDC} = 1$ means address $0xabc$ is supplying one USDC. Similarly for the distribution of borrows.

Step 2.5 Since some tokens might have very low usages during the stress test, leading to extremely high supply caps, we floor the borrow usage at 25% of borrowing power. To achieve this, we must first determine the amount of borrowed token collateralized by token A. Let $\vec{b}_i = \{b_{i,j}\}$ be the borrowed token balances for address i , and let $\beta_{i,A}$ denote the U.S. dollar amount of $\sum \vec{b}_i$ that is borrowed against token A in USD:

$$\beta_{i,A} = \frac{LT_A \cdot s_{i,A} \cdot p_A}{\sum_{j \in \text{tokens}} LT_j \cdot s_{i,j} \cdot p_j} \cdot \sum_{j \in \text{tokens}} b_{i,j} \cdot p_j \quad (3)$$

where LT_A is the liquidation threshold of A and p_j is the price of j . Notice that $LT_j \cdot s_{i,j} \cdot p_j$ denotes address i 's "borrowing power" with respect to token j , such that the fraction in Equation 3 denotes the fraction of their borrowing power that comes from a token A. It follows that $\beta_{i,A}$ denotes the fraction of address i 's borrows that are collateralized by a token A.

We ensure it is at least 25% of the borrowing power of token A:

$$\beta_{i,A} \geq 0.25 \cdot LT_A \cdot s_{i,A} \cdot p_A \quad (4)$$

meaning:

$$\frac{\sum_{j \in \text{tokens}} b_{i,j} \cdot p_j}{\sum_{j \in \text{tokens}} LT_j \cdot s_{i,j} \cdot p_j} \geq 0.25 \quad (5)$$

We achieve this by uniformly scaling up \vec{b}_i until it satisfies Equation 5.

Step 3 Scale the distribution of supplies and borrows by the supply cap, S_{cur} .

To scale the borrow vector of an address, we then distribute this scaling factor $\beta_{i,A}$ uniformly among this address's borrows. We know $\frac{\beta_{i,A}}{\sum \vec{b}_i}$ is the fraction of the borrows collateralized by token A so:

$$\vec{b}_i^* = \vec{b}_i \cdot \left(1 + \frac{\beta_{i,A}}{\sum b_{i,j} \cdot p_j} (\gamma - 1) \right) \quad (6)$$

where γ is our supply scaling factor:

$$\gamma = \frac{S_{cur}}{\sum \vec{s}} \quad (7)$$

Notice that $\gamma > 1$ if we increase the supply, $\gamma < 1$ if we decrease the supply. Furthermore, if $\gamma > 1$, the new borrow vector increases and vice-versa. To actually scale the supply, we scale each account's supply of token A:

$$s_{i,A}^* = s_{i,A} \cdot \gamma \quad (8)$$

Step 4 Initialize the agent-based simulation using \vec{s}^* and \vec{b}^* with the price trajectories defined in step 3. We arrive at an estimate of protocol losses under these conditions. Denote these losses as X .

Step 5 If $|X - K| < \varepsilon$, then stop, and set the supply cap S^* equal to S_{cur} . Otherwise, if $X > K$, then set $S_{max} = S_{cur}$, and set $S_{cur} = \frac{S_{min} + S_{max}}{2}$ with the new value of S_{cur} . Otherwise, if $X < K$, then set

$S_{min} = S_{cur}$, and set $S_{cur} = \frac{S_{min} + S_{max}}{2}$ with the new value of S_{cur} , and go back to step 3. This is the binary search step.

3.1.2 Assumptions

We have implicitly made three key assumptions in this methodology, described below.

1. **Elastic borrows and supplies:** Suppliers and borrowers are elastic relatives to the supply cap. If the supply cap is raised, then the supply of that token is raised accordingly. Similarly, the borrows against that token also rise. This relationship is encapsulated in Equations 6 and 8.
2. **Uniform borrower elasticity:** The increase in borrows for any address is distributed uniformly along all the borrowed tokens.
3. **Static usage:** Using a token as collateral stays constant as we scale the supply cap. This means that if 50% of an address' supply is being used to borrow, then when we scale the supplies and borrows, we retain this 50% usage.

We could repeat the supply cap experiment over multiple randomly sampled blocks to mitigate the effect of the above assumptions. Since the supply and borrow caps are snapshots of the block being examined, we repeat our stress tests with data from multiple blocks and average the resulting supply cap. This reduces the sensitivity of our results to the variance of any block's supply or borrows distribution. We recommend sampling blocks over multiple days (e.g. one block per day).

3.2 Choosing a Target Loss

The target loss K is the maximum amount of losses the protocol is willing to withstand. Since allowing users to supply and borrow the token is the primary function of Aave, there is a positive utility in increasing the supply cap. If our only goal was minimizing risk, we would set all supply caps to zero. Therefore, we establish a minimum capitalization ratio, similar to how traditional financial institutions impose risk capital requirements as a function of their risk-weighted assets. We could, for example, cap the protocol losses at $K = 0$ if we are looking to be conservative and there is plenty of on-chain liquidity for the token in consideration. Another potential candidate is some percentage of the safety module since the safety module is the ultimate source of insurance for suppliers. This paper considers an example target loss of 1% of the safety module, roughly \$4M as of March 2023. This example K is chosen for ease of exposition. Notice that the size of the safety module and the usage parameter described in the previous section all change over time. This requires risk providers to perform the stress tests described above consistently and frequently propose updates to risk parameters such as the supply cap.

3.3 New Token Listings

When listing new tokens on Aave, we must provide supply caps. However, we will not have the borrow and supply distribution required for our supply cap search. We consider two options:

1. Generate borrower and supplier vectors from some mathematical distribution, such as an exponential distribution, or a Dirichlet distribution.

2. Use the distribution of similar tokens. For example, if listing a new ETH liquid staking derivative, we might use the supplier and borrower distributions for stETH, cbETH, or rETH. We run the simulations using all three distributions and return the average supply cap.

As exponential and Dirichlet distributions would inevitably be fit against existing distributions (i.e., we don't know what a reasonable distribution of borrowers is unless we look at existing borrower distributions), we find that averaging results from existing distributions will likely yield better results. This is also unlikely to have a meaningful impact, as the liquidity risk primarily depends on the target token's on-chain liquidity and its corresponding stress test.

4 Discussion

Our methodology aims to mitigate the losses to the Aave protocol from foreseeable stress scenarios. We have described throughout the paper why stress testing is a valuable tool to inform us about protocol risks. Here, we consider some limitations in our methodology and potential future work.

4.1 Sensitivity Analysis

A crucial element of stress testing is sensitivity analysis: we consider how exaggerating the rate at which price and liquidity drops affects protocol losses. Additionally, this allows us to distinguish between riskier tokens at higher risk of depreciating faster than previously observed and safer tokens such as ETH. For example, the Baseline Token Drawdown scenario described in Section 2.1.4 tracks the price and liquidity of ETH during the Terra collapse. We also measure protocol losses if tokens observed a faster drawdown than ETH in price and liquidity. To achieve this, we modify Equations 1 and 2:

$$p_{A,t} = p_{A,0} \frac{p_{R,t}}{p_{R,0}} \cdot \lambda \quad (9)$$

$$L_{A,t} = L_{A,0} \cdot \frac{L_{R,t}}{L_{R,0}} \cdot \lambda \quad (10)$$

where λ is the sensitivity scaler, we consider scaling factors of (10%, 5%, 1%), leading to more conservative or more aggressive recommendations. This sensitivity analysis informs us of how sensitive protocol losses are to the liquidity of each token, allowing the community to make informed decisions on how conservatively to set the supply cap.

4.2 Additional Stress Tests

Our stress scenarios described in Section 2 are comprehensive but do not capture the idiosyncracies of particular asset classes. For example, we do not generate a scenario particular to ETH liquid staking derivatives, which might pose unique risks. Furthermore, we do not consider specific risks that Aave is subject to, for example, (1) a governance vote auctions off a token listed on Aave, leading to a sharp decline, or (2) a price manipulation attack against more illiquid tokens.

Developing these bespoke stress scenarios could provide invaluable insights into the protocol's risk management. However, they are left to future work. We provide stress scenario examples of governance auctions using MakerDAO's MKR as a reference event and an example of a short-seller attack using CRV as a reference event in Appendix B.

Furthermore, we only consider one reference event per stress scenario. We could instead consider various reference events and average their results to set the supply cap. However, the liquidity trajectory of older events might differ from that observed in more recent events. As we gather more reference events over longer time horizons, a potential improvement to our proposed methodology is to weigh each event by the time since it occurred. Weighing observations by the time since they have occurred is often done with exponentially weighted moving averages (EWMA) [4]. For example:

$$\begin{aligned}
 \alpha_i &\sim \text{Exp}(\lambda) \\
 \kappa &= \frac{1}{\sum \alpha_i} \\
 w_i &= \kappa \cdot \alpha_i \\
 \therefore \sum w_i &= 1
 \end{aligned}
 \tag{11}$$

where w_i are weights, α_i is i.i.d. from an exponential distribution, and κ ensures that our weights sum up to 1.

4.3 Synthetic Stress Tests

Another common stress testing framework for financial institutions is testing against “what-if” scenarios. These scenarios are manually generated to achieve a certain effect. For example, we might simulate protocol losses if a large borrower takes out a big position on a token, and it suddenly depreciates in value. This could allow us to simulate the consequences of a short-seller attack on the protocol.

The Federal Reserve, for instance, generates stress scenarios that are categorized by hypothetical macroeconomic trends. For the 2022 severely adverse scenario provided to bank hold companies, the Fed considered a global economic recession, an acceleration in remote work, a sharp increase in unemployment, etc. [5]. Here we have proposed stress tests based exclusively on historical data on the referenced events. We could extend this methodology to consider other potential stress tests based on wider macroeconomic data. For instance: testing various tokens against the scenario where BTC suddenly halves price. Generating reasonable price and liquidity trajectories for these synthetic scenarios is left to future work.

4.4 Hollistic Stress Testing

Our stress testing is not holistic: we don’t measure protocol losses if all tokens are stressed simultaneously. For this reason, we might want to work on developing a more holistic measure of risk, where we simulate many classes of tokens being stress tested at once. Then we see how to mitigate this risk by setting some risk parameters. This would greatly increase the complexity of our program, and we leave it for future work.

4.5 Applications to Other Protocol Risk Parameters

We can extend this methodology with minor adjustments to simulate the losses due to isolated assets. We then apply a similar binary search to set the debt ceilings of isolated assets, with minor adjustments to the methodology.

A Supply Cap Search Algorithm

Algorithm 1 Supply Cap Search Algorithm.

```
1:  $S_{min} \leftarrow 0$ 
2:  $S_{max} \leftarrow$  (current circulating supply of A)
3:  $S_{cur} \leftarrow \frac{S_{min} + S_{max}}{2}$ 
4:  $i \leftarrow 0$  ▷ Accumulator
5: while  $i <$  (maximum number of iterations) do
6:    $\vec{s} \leftarrow$  (the current supply vector for all Aave users)
7:    $\vec{b} \leftarrow$  (the current borrow vector for all Aave users)
8:   Scale  $\vec{s}$  according to Equation 8
9:   Set  $\vec{b}$  according to Equation 6
10:
11:   Set price trajectory,  $p_{A,t}$ , according to Equation 1
12:   Set liquidity trajectory,  $L_{A,t}$ , according to Equation 2
13:    $X \leftarrow$  (simulated losses with  $p_{A,t}$  and  $L_{A,t}$ )
14:   if  $|X - K| \leq \varepsilon$  then
15:     return  $S^* \leftarrow S_{cur}$ 
16:   else
17:     if  $X > K$  then
18:        $S_{max} \leftarrow S_{cur}$ 
19:     else
20:        $S_{min} \leftarrow S_{cur}$ 
21:     end if
22:      $S_{cur} \leftarrow \frac{S_{min} + S_{max}}{2}$ 
23:   end if
24:    $i \leftarrow i + 1$ 
25: end while
26: return  $S^* \leftarrow S_{cur}$  ▷ our best estimate of the supply cap
```

B Additional Stress Scenarios

Reference Event 1: MKR’s black Thursday. MakerDAO auctioned off a significant portion of MKR’s circulating supply to cover losses from missed liquidations in March of 2020. The auction of roughly 21 thousand MKR tokens incurred slippage of around 50%, leading the token to trade at \approx \$270 USD instead of the \approx \$500USD it was previously trading at.

- Start Time: March 11, 2020, 12:00:00 AM
- End Time: March 14, 2020, 12:00:00 AM

Project	Version	Contract	USD Vol	Pct Vol
Uniswap	1	0x2c4bd064b998838076fa341a83d007fc2fa50957	6,090,879	76.27%
Kyber	1	0x65bf64ff5f51272f729bdcd7acfb00677ced86cd	1,646,514	20.62%
Oasis	2	0x794e6e91555438afc3ccf1c5076a74f42133d08d	199,242	2.49%

Table 6: MKR DEX Liquidity during the MKR reference event.

Reference Event 2: CRV’s economic exploit. The Curve DAO token, CRV, was exploited on November 22nd, 2022.

- Start Time: November 12, 2022, 12:00:00 AM
- End Time: November 24, 2022, 12:00:00 AM

Project	Version	Contract	USD Vol	Pct Vol
Curve	Regular	0x8301ae4fc9c624d1d396cbdaa1ed877821d7c511	163,771,609	60.72%
Curve	Factory V1 Plain	0x9d0464996170c6b9e75eed71c68b99ddedf279e8	30,734,595	11.39%
Uniswap	3	0x4c83a7f819a5c37d64b4c5a2f8238ea082fa1f4e	17,229,456	6.39%
Uniswap	3	0x919fa96e88d67499339577fa202345436bcdaf79	14,979,039	5.55%
1inch Limit Order Protocol	2	0x119c71d3bbac22029622cbaec24854d3d32d2828	12,461,967	4.62%
1inch Limit Order Protocol	eRFQ v1	0x1111111254fb6c44bac0bed2854e76f90643097d	8,164,407	3.03%
Sushiswap	1	0x58dc5a51fe44589beb22e8ce67720b5bc5378009	7,262,149	2.69%
Uniswap	3	0x9445bd19767f73dcae6f2de90e6cd31192f62589	4,954,000	1.84%
0x Native	4	0xdef1c0ded9bec7f1a1670819833240f027b25eff	4,934,789	1.83%
Uniswap	2	0x3da1313ae46132a397d90d95b1424a9a7e3e0fce	3,635,085	1.35%

Table 7: CRV DEX Liquidity during the LUNA reference event.

References

- [1] Martin Cihak. “Introduction to applied stress testing”. In: *SSRN Electronic Journal* (2007). DOI: [10.2139/ssrn.973989](https://doi.org/10.2139/ssrn.973989).
- [2] Omer Goldberg and Yonatan Haimovich. “Aave V3 Risk Parameter Methodology”. In: (2023).
- [3] Yonatan Haimovich and Omer Goldberg. “A Note on Isolation Mode Risk Parameters”. In: (2023). URL: https://chaoslabs.xyz/resources/chaos_aave_isolation_mode.pdf.
- [4] John Hull. “Chapter 22”. In: *Risk management and financial institutions*. Wiley, 2023.
- [5] Federal Reserve. *Dodd-Frank Act Stress Test Publications*. URL: <https://www.federalreserve.gov/publications/2022-Stress-Test-Scenarios.htm>. accessed: 03/20/2023.