CHAOS LABS

Towards a Quantitative Interest Rates Framework

An initial exploration of optimizing interest rates through comprehensive analysis of supply-demand dynamics and careful adjustments for asset risk.

AAVE

April 14, 2023

Interest Rates Optimization Framework Exploration

Omer Goldberg
omer@chaoslabs.xyz

Yonatan Haimowitz
haimo@chaoslabs.xyz

Shai Kritz
shai@chaoslabs.xyz

April 2023

Abstract

This study undertakes a comprehensive exploration of the Aave protocol, an open, multi-chain borrowing-lending mechanism that allows participants to motionlessly borrow and lend crypto assets. The interest rates for both lenders and borrowers are determined by the Interest Rate Strategies, which have largely remained unaltered since the protocol's initiation.

In this analysis, we delve into the supply and demand dynamics within Aave and introduce an innovative methodology for revising interest rate parameters. Initially, we elaborate on the notion of optimal utilization and propose a strategy based on the token's risk collateral. Subsequently, we justify the need to aim for optimal utilization by modulating the pool's interest rate parameters. The concluding segment of this study explores the plausible outcomes steered by our proposed approach.

Contents

Disclaimer

This document is purely informational and does not constitute an invitation to acquire any security, an appeal for any purchase or sale, or an endorsement of any financial instrument. Neither is it an assertion of the provision of investment consultancy or other services by Chaos Labs Inc. References to specific securities should not be perceived as recommendations for any transaction, including buying, selling, or retaining any such securities. Nothing herein should be regarded as a solicitation or offer to negotiate any security, future, option, or other financial instrument or to extend any investment advice or service to any entity in any jurisdiction. The contents of this document should not be interpreted as offering investment advice or presenting any opinion on the viability of any security, and any advice to purchase, dispose of or maintain any security in this report should not be acted upon. The information contained in this document should not form the basis for making investment decisions.

While preparing the information presented in this report, we have not considered individual investors' specific investment requirements, objectives, and financial situations. This information does not account for the specific investment goals, financial status, and individual requirements of the recipient of this information, and the investments discussed may not be suitable for all investors. Any views presented in this report by us were prepared based on the information available when these views were written. Additional or modified information could cause these views to change. All information is subject to possible rectification. Information may rapidly become unreliable for various reasons, including market or economic changes.

Chapter 1

Overview

1 Motivation

The Aave protocol, a key player in decentralized finance (DeFi), significantly influences interest rate curves for a variety of tokens through its pre-established strategies. These strategies, occasionally refined by community contributions and service providers, were essential during the protocol's early stages, when data was sparse, and risk management tools were undeveloped. As Aave continues to evolve, we propose a framework to enhance these interest rate strategies by utilizing data on equilibrium borrow and supply rates and identifying when parameter adjustments are required.

Market borrowing and external factors predominantly influence supply rates and can exhibit volatile trends. An exemplary case is the impact of Lido's staking yield on the market rates for WETH on Aave. To address these inherent challenges, we propose a systematic approach that enables regular, modest adjustments to interest rate curve parameters, focusing specifically on Month over Month (MoM) changes in observed borrow and supply rate equilibriums.

This study introduces an initial iterative method for configuring interest rate parameters, relying heavily on empirical borrower and lender behavior data. Our primary objective is to optimize utilization by adjusting interest rate parameters, thereby boosting capital efficiency for suppliers and augmenting the protocol's reserves while promoting increased borrowing activity.

The study is divided into four sections. The first section establishes the means to determine optimal utilization, while the second advocates targeting this optimal utilization using Aave's interest rate parameters. The third section proposes a set of parameter modifications based on January to April 2023 data and scrutinizes their potential effects. The final section explores the empirical elasticity of borrowers and suppliers. We propose several statistical methodologies to assess elasticity but conclude that their statistical significance is insufficient to serve as our primary reference. Instead, we evaluate potential outcomes based on relative borrower and supplier elasticity, employing graphical representations and statistical tests to guide readers toward the most probable outcomes.

Chapter 2

Interest Rate Curves Primer

1 Introduction

Interest rate curves in Aave, as well as most other lend-borrow protocols, are fundamentally derived from the usage, or ***utilization***, of a particular pool:

$$u = \frac{\text{Amount Borrowed}}{\text{Amount Supplied}} \tag{2.1}$$

Utilization directly influences the interest rate applied to borrowings within the protocol, establishing a linear correlation. As borrowing demand surges or supply diminishes, both utilization and borrow rates subsequently escalate. In parallel with margin trading protocols like Compound and Euler, Aave delineates the interest rate for a pool using a jump-rate model. Herein, upon reaching a certain level of utilization, denoted as the *optimal utilization* u_{opt}, the slope of the interest rate curve sharply increases:

$$r_b = f_b(u) = \left\{ \begin{array}{r} r_0 + r_1 \frac{u}{u_{opt}}, \\ u \leq u_{opt} \\ r_0 + r_1 + r_2 \frac{u - u_{opt}}{1 - u_{opt}}, \\ u > u_{opt} \end{array} \right\} \tag{2.2}$$

$$r1 << r2 \tag{2.3}$$

Where r_b is the interest rate charged to borrowers, and r_0, r_1, r_2 are all constants we must set. The interest rebated to suppliers is derived from the borrowing rate:

$$r_s(u) = f_s(u) = \begin{Bmatrix} u \cdot \left(r_0 + r_1 \frac{u}{u_{opt}}\right)(1 - RF), \\ u \leq u_{opt} \\ u \cdot \left(r_0 + r_1 + r_2 \frac{u - u_{opt}}{1 - u_{opt}}\right)(1 - RF), \\ u > u_{opt} \end{Bmatrix} \quad (2.4)$$

RF is the reserve factor, and the interest portion accrues to the protocol's reserves. We ignore stable borrowing for now.

1.1 WETH Investigation

The WETH market on Aave v2 provides abundant empirical data regarding borrower and supplier behavior. Understanding these behavioral patterns is crucial for our methodology in setting interest rates, specifically how borrowers and suppliers respond to their respective borrow and supply rates, subsequently affecting utilization. As we argue in this paper, regular and thoughtful modifications to interest rate parameters can help build statistical confidence concerning borrower and supplier elasticity.

Since January 2022, there have been three adjustments to interest rate parameters. In each case, the alterations were primarily motivated by the aim to align Lido's staking rate with the borrowing rate at optimal utilization. Incorporating stETH into Aave v2 in late February 2022 promptly led to a surge in utilization.

March 2022

A proposal was put forward in March 2022, suggesting a substantial reduction of the ETH borrow rate to stimulate heightened utilization. This amendment instantaneously diminished borrowing rates for users, promoting further borrowing:

$$r_{slope1} : 8\% \rightarrow 3\%, \;\; u_{opt} : 65\% \rightarrow 70\% \quad (2.5)$$

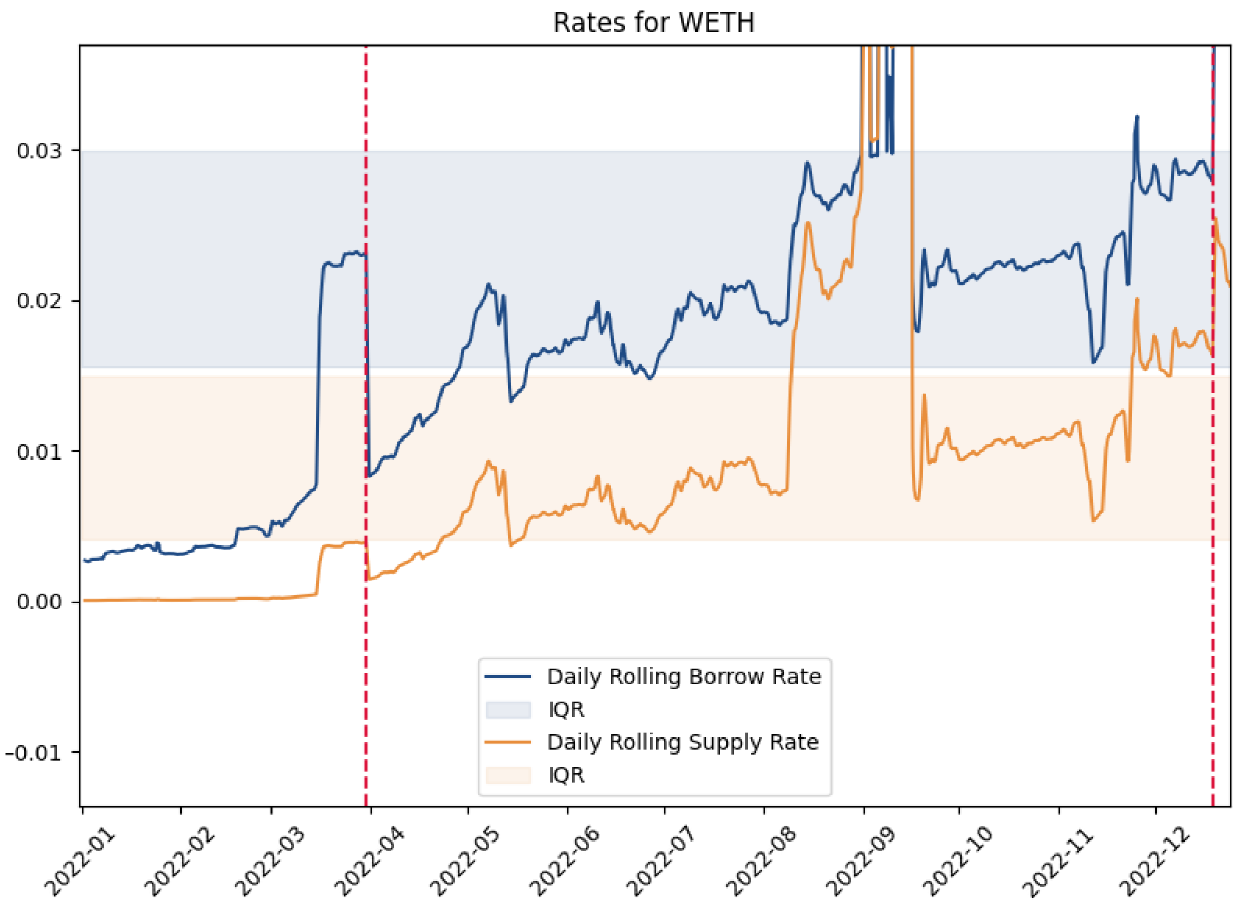

Figure 2.1: Red dotted lines indicate parameter changes. Notice how on the first change in March 2022, borrow rates instantly drop from 2.5% to 1%, increasing over time to track Lido's stake rate. This mean-reversion from the deflated rates following the March 2022 proposal to the underlying market rate of 3% suggests borrower elasticity.

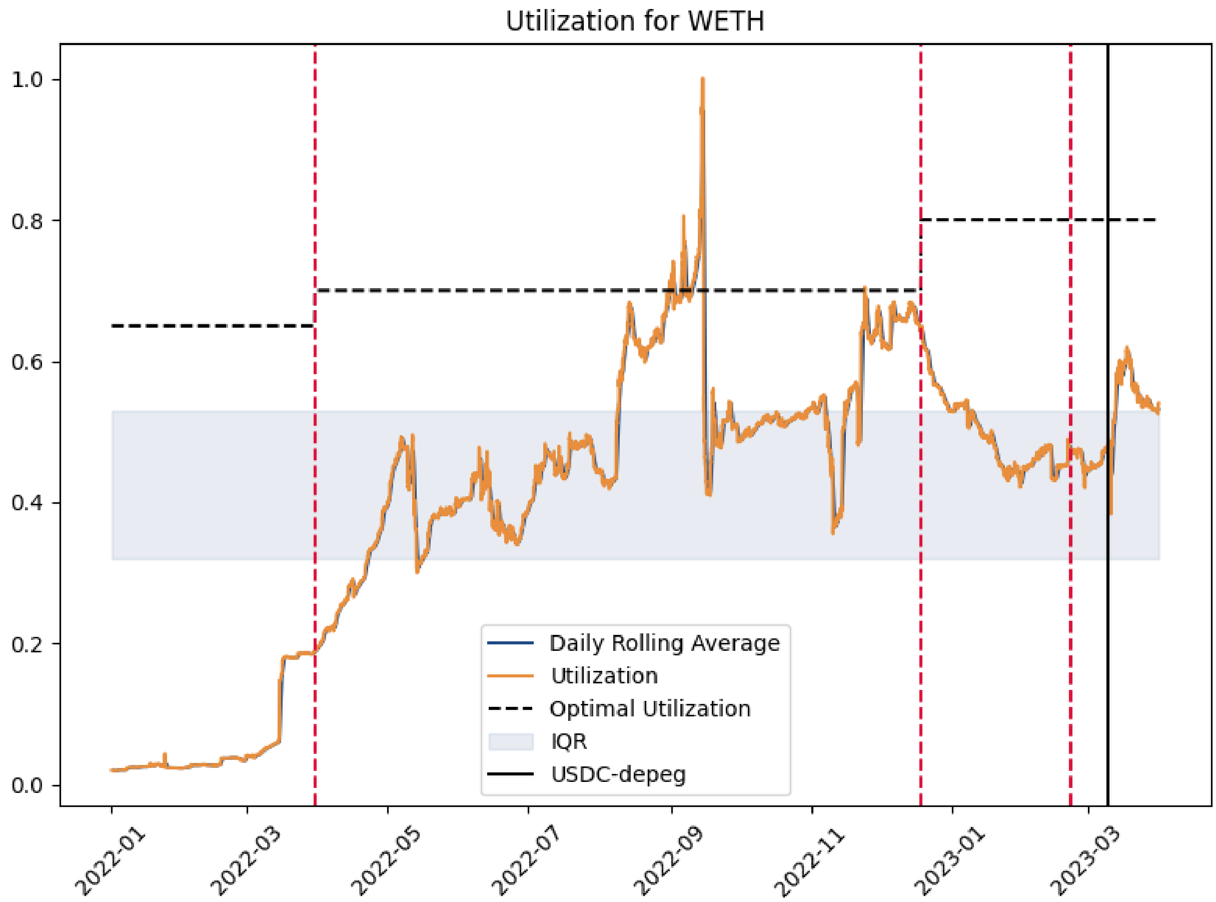

Figure 2.2: Here utilization increases significantly throughout 2022 as Lido deposits become eligible as collateral on Aave, allowing users to stake WETH and borrow against their wstETH recursively. This new source of yield encourages users to borrow at higher rates, meaning higher utilization. With the IR curve being flattened by the proposal on March 2022, this effect is magnified, and utilization massively increases from 1% to 50% towards the end of 2022 and early 2023.

Upon activating stETH collateral, borrowers demonstrated elasticity in response to the borrowing rate on Aave. By reducing r_{slope1}, Aave successfully boosted utilization, noting an increase in borrowing that effectively returned the borrowing rate to the market rate. The "market rate" for borrowing on Aave is equivalent to the staking yield on Lido: as staking yields increase, so does the rate borrowers are willing to pay to borrow WETH and stake it. As a result, it can be inferred that the "market rate" for Aave is largely exogenous, necessitating Aave's service providers to adjust interest rate parameters regularly to target optimal utilizations.

This adjustment highlights a key feature of interest rate parameters on Aave. Although the community doesn't have control over market borrow or supply rates, a thorough understanding of borrower and supplier elasticity could allow for control over the point of market equilibrium. For instance, Lido's yield determines how much borrowers are willing to pay (for recursive staking) and what suppliers will charge (due to the opportunity cost of Lido staking). However, the Aave community can influence whether these market rates occur at 10% utilization or a higher level of 60% utilization. Assuming borrowers or suppliers exhibit elasticity in response to interest rates, which isn't a given, they will react to changes in rate parameters.

	symbol	start	end	supply is elastic	supply elasticity stat	lag	demand is elastic	demand elasticity stat	lag
0	WETH	2022-01-01 00:00:10	2022-04-01 00:00:10	False	1.332	1	False	0.842	1
0	WETH	2022-04-01 00:00:10	2022-06-30 00:00:10	False	1.344	24	True	4.937	24
0	WETH	2022-06-30 00:00:10	2022-09-28 00:00:10	True	3.455	5	True	28.566	5
0	WETH	2022-09-28 00:00:10	2022-12-27 00:00:10	True	2.927	9	False	0.257	9
0	WETH	2022-12-27 00:00:10	2023-03-27 00:00:10	True	2.866	13	True	9.037	13

Figure 2.3: Using the output of Granger-causality tests conducted on three-month intervals, we can substantiate, to a certain degree, our hypothesis that borrowers and suppliers become elastic to interest rates as recursive LSD borrowing is enabled on Aave. The periods from January 2022 to April 2022, September 2022 to December 2022, and January 2023 to March 2023 present noisy data: structural breakpoints (i.e., the interest rate changes) render autoregressive statistical tests such as Granger-causality tests less informative. Nonetheless, we can glean some evidence that the supply of WETH, particularly the demand for WETH, demonstrates elasticity concerning interest rates.

In the following sections, we will discuss statistical tests that provide evidence supporting our theory that the supply and demand for WETH are elastic. Our results indicate that when stETH is listed on Aave and borrow rates are reduced, supply and demand demonstrate more statistically significant elasticity.

1.2 February 2023

A consequential modification to the parameters transpired in 2023, providing insights that significantly inform our methodology. Following the enactment of AIP 131 in late 2022, the borrowing activity associated with WETH had stagnated and even diminished, consequently leading to a decrease in the accrual of Aave's reserves. To stimulate more WETH borrowing, service providers suggested a decrease in r_{slope1} and an increase in the base rate r_0. Regrettably, these adjustments produced conflicting effects at the pertinent utilization levels and did not result in any substantial enhancement in utilization:

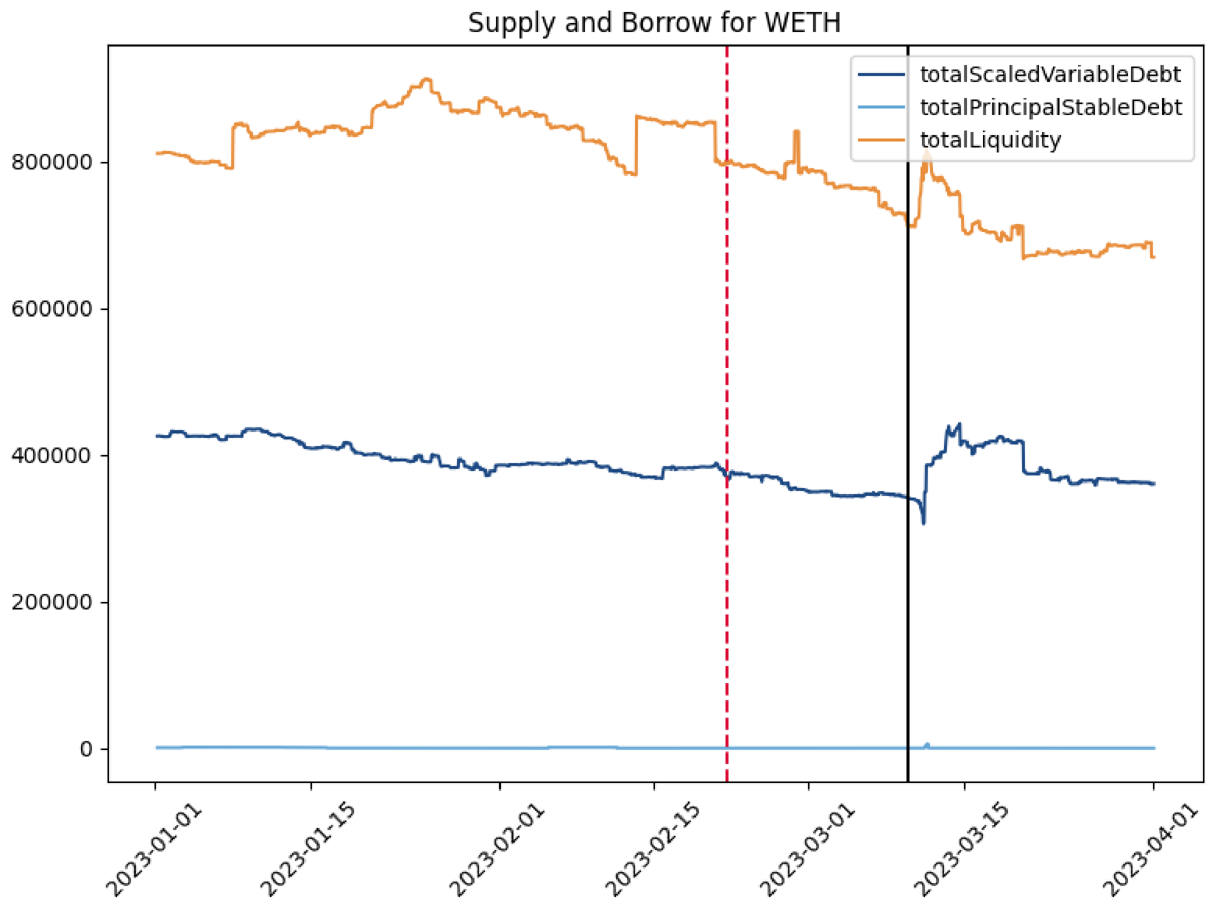

Figure 2.4: The orange line (supply) and the blue line (demand) before and after the proposal by Llama on February 2023 (dotted line). Notice that the borrowing activity did not meaningfully change after implementing the change. The black line indicates the USDC depeg, which created noise in the data on various markets.

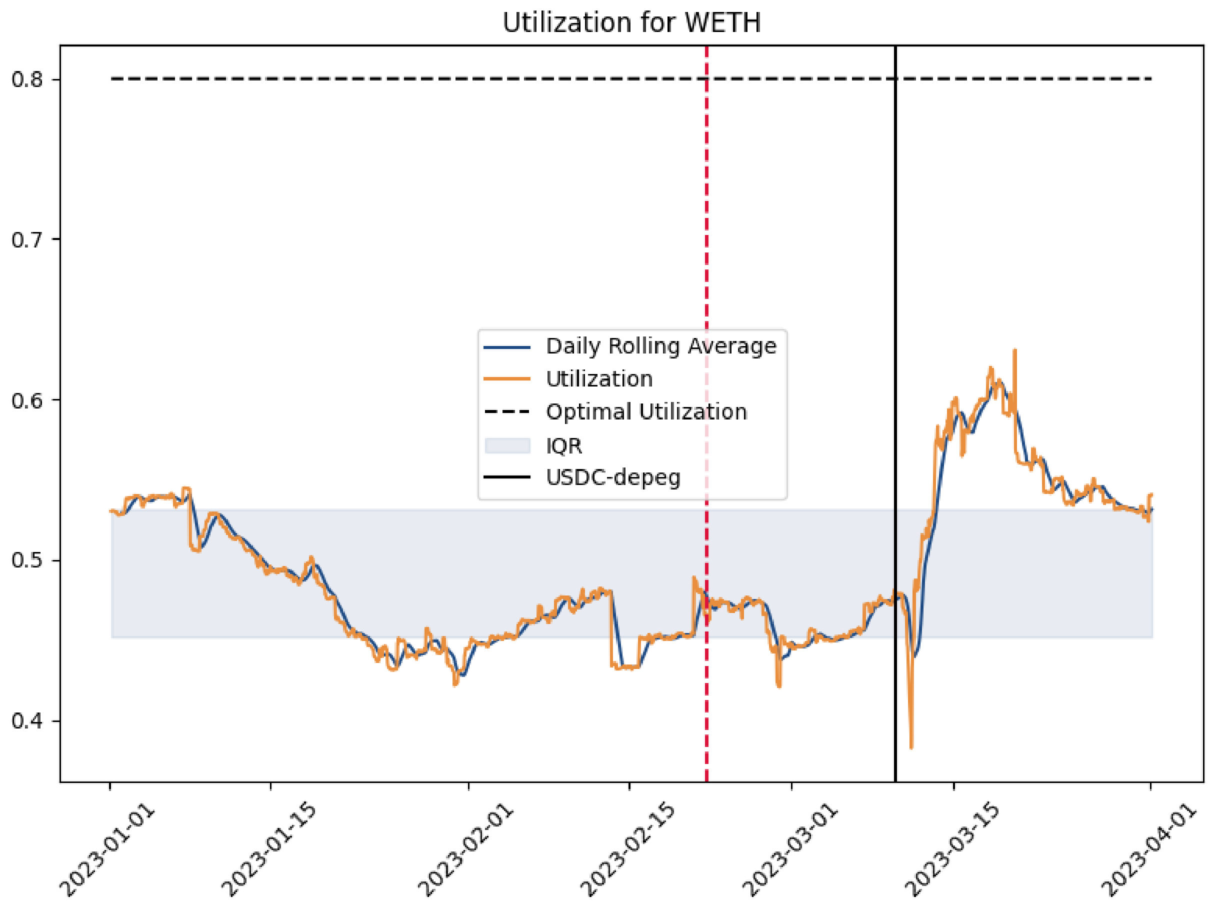

Figure 2.5: Notice that the utilization remained largely contained around the 50% mark despite the material reduction in slope 1 from 5.75% to 3.8%.

The proposed changes were unlikely to improve utilization on the pool due to the increase in the base rate for the pool:

$$r_{\text{slope1}} : 5.75\% \rightarrow 3.8\%, r_0 : 0\% \rightarrow 1\% \tag{2.6}$$

The proposed alteration resulted in a flattening of the borrow curve around the current utilization. Consequently, neither borrowers nor suppliers noted a significant change in their respective rates pre and post-amendment. Although this could potentially diminish interest rate volatility, which might be advantageous for both borrowers and suppliers, it does not induce any behavioral modification. In the results section, we will recommend a gradual reduction of the base slope to encourage an uptick in utilization.

As we compile more data pertaining to the impact of various interest rate-related AIPs, we may apply structural breakpoint tests, such as the Chow test or a Bayesian structural breakpoint test, to ascertain whether these AIPs induce statistically significant shifts in the underlying supply and demand trends.

Chapter 3

Optimal Utilization

1

Here we overview what optimal utilization is, why it is *optimal*, and how we can set it using sophisticated risk metrics.

1.1 Why is the Optimal Utilization "Optimal"?

To comprehend our definition of *optimal* utilization, it's crucial to consider the collective welfare of the protocol participants. Typically, four types of agents are involved in the protocol: borrowers, suppliers, liquidators, and the protocol itself. Initially, let's focus on borrowers and suppliers.

The utility functions of borrowers and suppliers are in contrast. Borrowers aspire to borrow at lower rates, while suppliers aim to earn at higher rates. Let's assume that borrowers and suppliers equally contribute to our total welfare calculation: our goal is to maximize the utility of both borrowers and suppliers on an equal footing. Essentially, this implies our intention to minimize the spread between borrow and supply rates:

$$S(u) = \left\{ \begin{array}{r} \left(r_0 + r_1 \frac{u}{u_{opt}}\right)(u \cdot (1 - RF) - 1)\,, \\ u \leq u_{opt} \\ \left(r_0 + r_1 + r_2 \frac{u - u_{opt}}{1 - u_{opt}}\right)(u \cdot (1 - RF) - 1)), \\ u > u_{opt} \end{array} \right\} \tag{3.1}$$

If we consider this spread S as a percentage of the borrowing rate r_b we find a linearly increasing function:

$$\frac{S(u)}{r_b} = (u \cdot (1 - RF) - 1) \tag{3.2}$$

This is intuitive; higher utilization implies higher capital efficiency, which is generally suitable for all parties involved. Naively minimizing this spread would lead us to believe that 100% utilization is optimal.

However, a 100% utilization rate might deter liquidators from handling underwater positions due to the unavailability of the underlying tokens as rewards once the liquidator settles the debt from the underlying positions. The liquidator can only claim the aToken representation of that collateral, which obliges them to bear **inventory risk** until the utilization drops below 100%, enabling the liquidator to exchange their aToken for the token itself. Thus, 100% utilization poses a challenge for the protocol due to the risk of potential liquidation omissions. Moreover, at 100% utilization, suppliers can not withdraw their funds, adversely affecting the protocol's users and culminating in an unfavorable user experience.

Consequently, the *optimal utilization* can be more accurately perceived as the *highest* utilization level that the protocol is prepared to risk, to avert the possibility of missed liquidations. The non-utilized supply serves as risk capital for the protocol: a safeguard of funds we must preserve for liquidations and withdrawals.

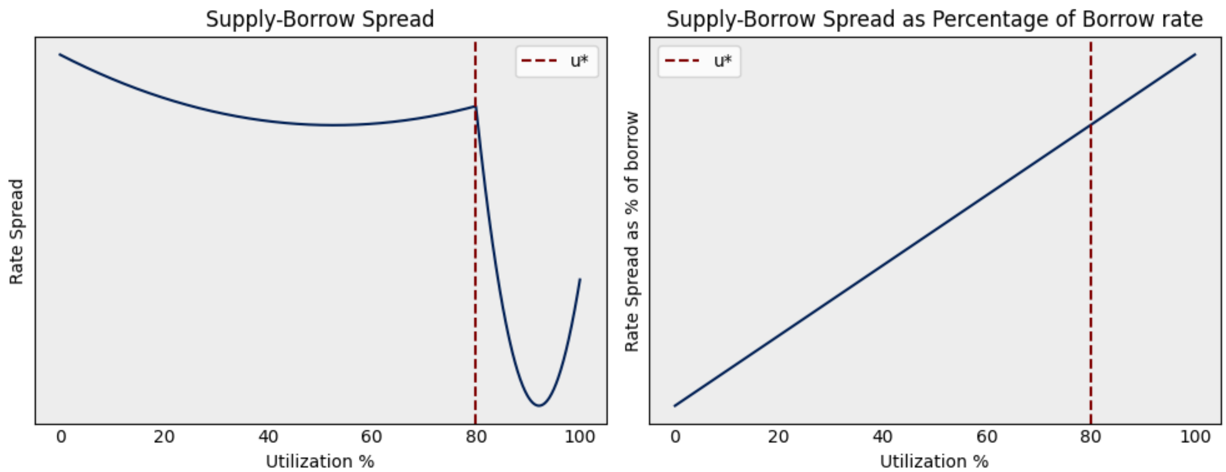

Figure 3.1: Notice that the utilization remained largely contained around the 50% mark despite the material reduction in slope 1 from 5.75% to 3.8%.

The risk of utilization approaching 100% will be denoted herein as *utilization risk*.

1.2 Determining Optimal Utilization

We've previously discussed several reasons why the optimal utilization must fall short of 100%. Our primary concern as risk managers for the protocol is ensuring sufficient risk capital to execute liquidations during abrupt asset price volatility episodes. We propose determining the optimal utilization for any specific pool to ensure adequate funds are available to liquidate all *at risk* positions. The amount that must be readily available for liquidations is the collateral at risk. We calculate this collateral at risk using the Chaos Labs' Collateral at Risk (CCAR) platform. Using the Collateral at Risk (CaR) definition we set:

$$u_{\text{opt}} = 1 - \frac{\text{CaR} + \mathbb{I}(.)\sigma}{\text{Supply}} \tag{3.3}$$

where $\mathbb{I}(.) = 0$ for a **neutral** recommendation, $\mathbb{I}(.) = 1$ for a **conservative** recommendation, and $\mathbb{I}(.) = -1$ for an **aggressive** recommendation. The indicator function aids us in tempering our recommendations according to the volatility of the asset's collateral at risk and, crucially, empowers Aave community members with a decision based on the community's risk tolerance and the strategic significance of any given token to the overall well-being of the protocol.

1.3 Availability of Collateral at Risk Data for Certain Tokens

Specific tokens might not possess data pertaining to collateral at risk as they may not have been approved as collateral. Regardless, elevated utilization leads to a situation where suppliers cannot retract their token supply. We propose the following framework for discerning optimal utilization.

We introduce a concept called **supply drawdown**, defined as a percentage decrease in supply over a designated period. For example, the maximum one-hour drawdown signifies the most pronounced percentage drawdown from the peak to the trough during a one-hour interval. As evidenced by multiple service providers on Aave, pool utilization rarely remain above optimal for durations surpassing approximately twelve hours. To adopt a prudent approach, we propose establishing optimal utilization by considering drawdown periods spanning from one hour up to twenty-four hours:

$$u_{\text{opt}} = 1 - (\text{p-Drawdown}) \tag{3.4}$$

Where p-drawdown indicates the pth percentile drawdown. For example, the 0th percentile drawdown is the highest drawdown observed. Intuitively, we ensure that net supply outflows lasting up to twenty-four hours can be covered before reaching 100%

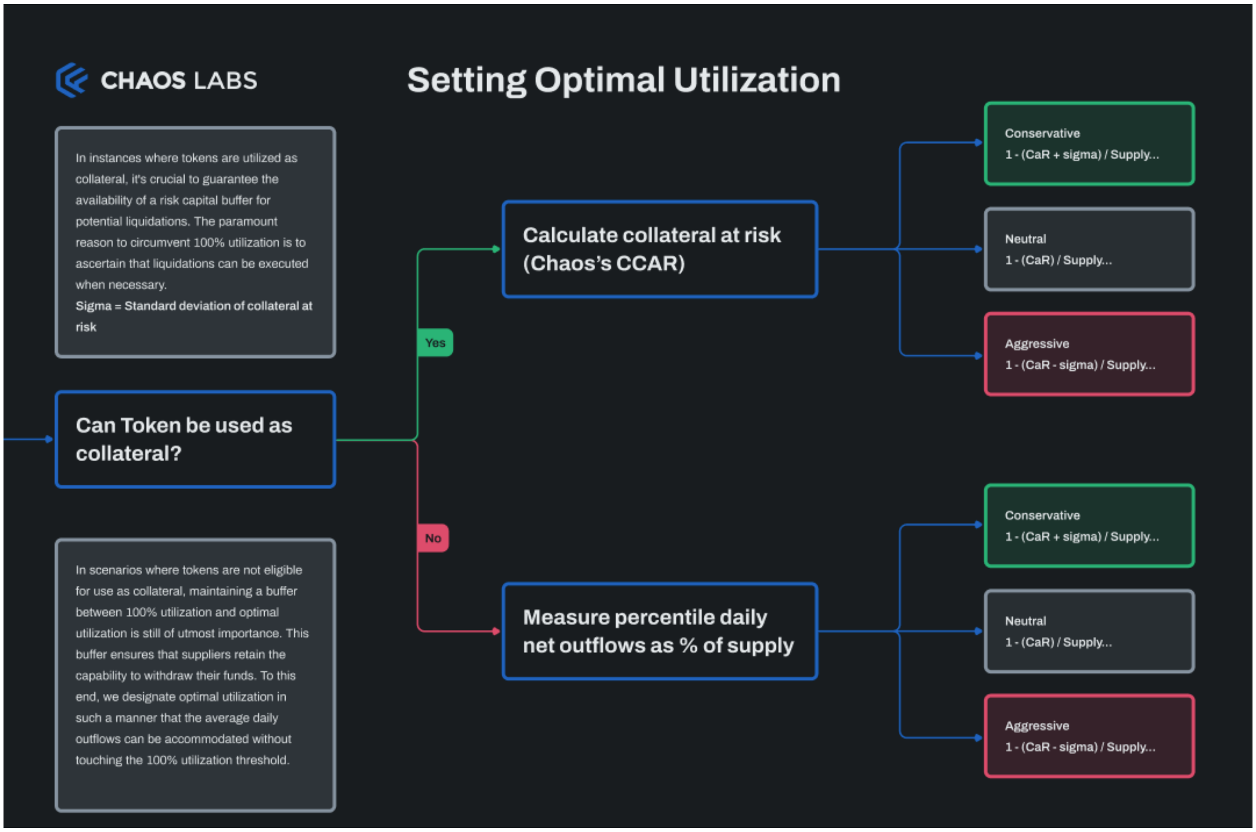

Figure 3.2: Summary: The following flow chart illustrates the decision tree used to determine the optimal utilization.

2 Targeting Optimal Utilization

To effectively target optimal utilization, we first must understand the dynamics of equilibrium interest rates on Aave. In traditional finance, equilibrium is typically represented by the market-clearing interest rate, the rate at which supply of funds matches the demand, leaving no excess supply or demand. In the context of Aave, this would suggest an equilibrium where utilization reaches 100%, with borrow and supply rates potentially soaring as high as 100%-300%. However, both empirical data and intuitive reasoning indicate that the balance of supply and demand in these lending protocols does not occur when supply equals demand, but rather when there is a surplus of supply.

To better understand how equilibrium rates form on Aave, let's define the **equilibrium supply rate** and the **equilibrium borrow rate**. Assuming there have been no changes to the interest rate parameters, a liquid borrowing/lending market will settle towards an equilibrium where borrowers are willing to accept the borrowing rate and suppliers are willing to provide funds at the supply rate. If a supplier withdraws from the market, the utilization increases (thus increasing both the borrow and supply rates), triggering either borrowers' repayments or additional deposits from suppliers, bringing the market back to equilibrium. In this context, we're implicitly assuming that either borrowers, suppliers, or both are **elastic** with respect to the current equilibrium borrow and supply rates.

The market equilibrium is disrupted if changes are made to the interest rate parameters. The current equilibrium supply rate would likely occur at a different utilization than the market supply rate. As a result, a new equilibrium would form at a different utilization level.

For example, let's consider a pool with the following parameters:

$$
\begin{aligned}
r_0 &= 0 \\
r_1 &= 0.05 \\
r2 &= 1 \\
RF &= 0.1 \\
u_{\text{opt}} &= 0.8
\end{aligned}
$$

The current equilibrium has remained at around $u^* = 0.5$, meaning the equilibrium borrow rate is $r_b^* \approx 0.031$, and the equilibrium supply rate occurs at $r_s^* \approx 0.014$. Suppose we aim to increase utilization by lowering the slope $r_1 \to 0.04$. The current equilibrium borrow rate now occurs at:

$$u^*_{\text{borrow}} = 0.625 = f_s^{-1}(r_b^*) \tag{3.5}$$

$$u^*_{\text{supply}} = 0.559 = f_b^{-1}(r_s^*) \tag{3.6}$$

To forecast how our change in r_1 will affect the pool's utilization, we must have some conception of whether borrowers and suppliers are elastic and if they are, who is more elastic.

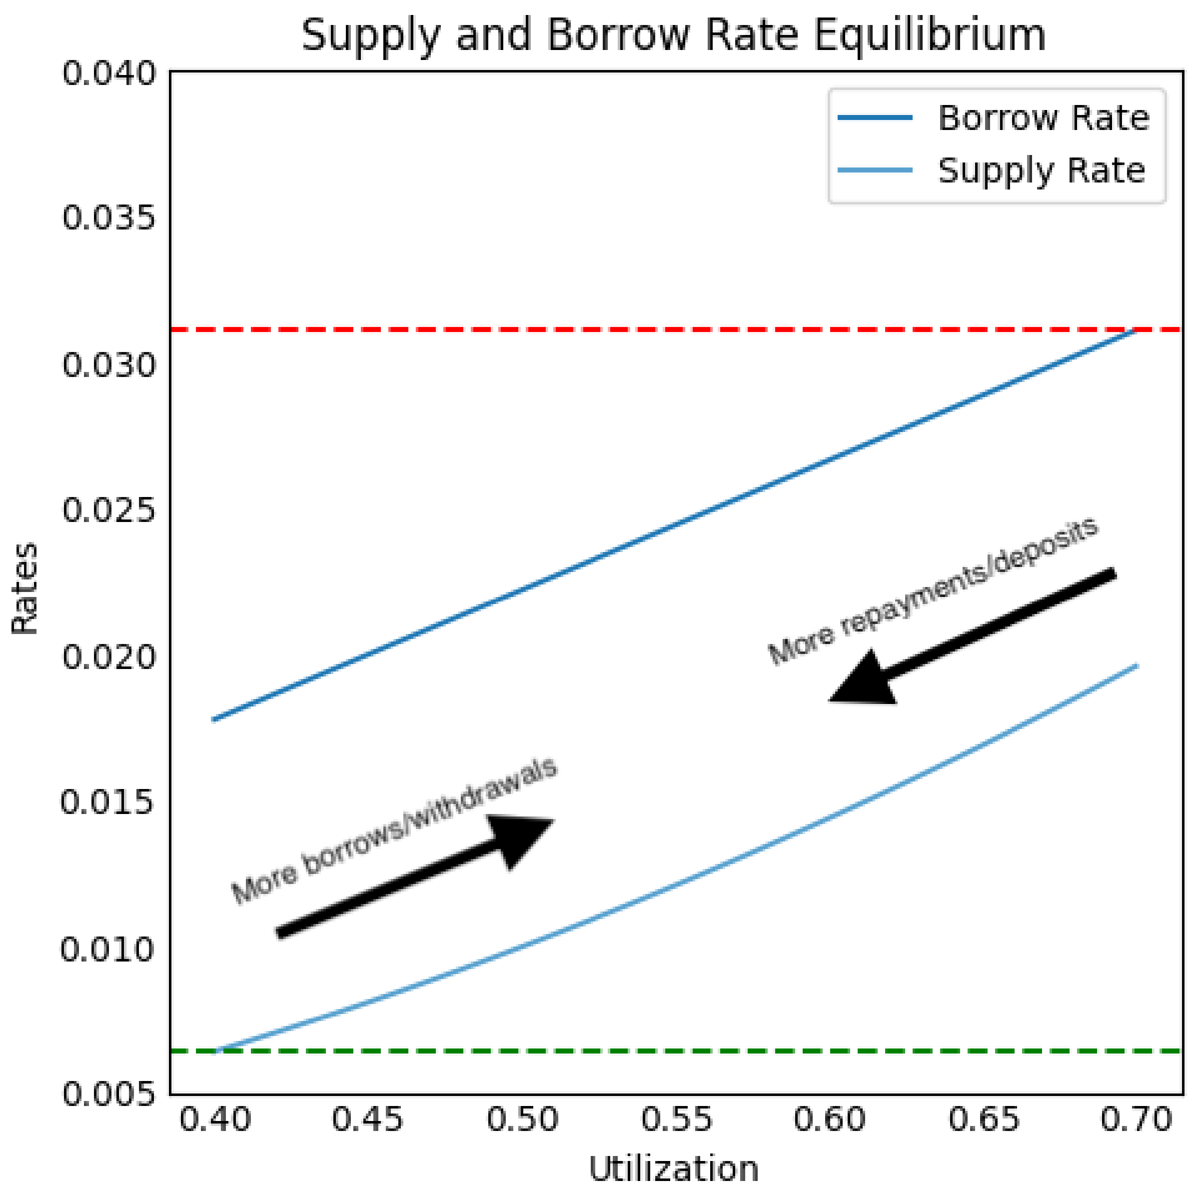

Figure 3.3: The green dotted line indicates the previous equilibrium supply rate, whereas the red dotted line indicates the previous equilibrium borrow rate. To understand what happens to utilization, we must know whether borrowers or suppliers are elastic (and therefore whether we will see changes in borrow/supply activity) and who is more elastic.

We have some confidence that the equilibrium utilization would occur somewhere between:

$$u^*_{\min} = f_s^{-1}(r_s^*)$$
$$u^*_{\max} = f_b^{-1}(r_b^*)$$

which in the figure corresponds to the range $[0.4, 0.7]$. The equilibrium within this range relies on the relative elasticity of borrowers and suppliers. If suppliers are significantly more elastic than borrowers, we'd observe an equilibrium closer to 0.4, and vice versa.

Naturally, there isn't a defined equilibrium supply rate or equilibrium borrow rate, but these terms serve as guiding principles to comprehend how we can target optimal utilization. We derive these equilibrium rates from recent supply and borrow rates, and consider probable outcomes within the inter-quartile range of the recently observed rates.

2.1 Regime Switches

Equilibrium rates may fluctuate due to endogenous factors, such as changes in interest rate parameters, and exogenous factors, such as alterations to staking yields. As we accumulate more data on changes in interest rate parameters, like the example with WETH in Section 1, statistical tests like the Chow test or Bayesian structural breakpoint tests may assist us in determining if there was a regime shift in equilibrium rates or utilization. This not only aids us in understanding the trends of supply and borrow rates but also whether borrowers and suppliers are elastic to parameter changes in the first place.

This paper proposes a simple and transparent method for guiding regular and moderate parameter adjustments in the absence of parameter change data.

We occasionally refer to the protocol's reserve accrual as "revenue". Reserves represent the protocol's only "revenue generation" mechanism, serving two potential purposes. As mentioned in the docs, the protocol's reserves (both from the reserve factor and the Aave ecosystem reserves) are utilized to sustain the DAO and compensate protocol contributors. On the other hand, the safety module serves as the primary source of insolvency protection. The reserve collector contract for Ethereum deployments can be found here, with approximately \$17M USD in reserves as of April 12th, 2023. Although the reserves contract has not yet been utilized, it will inevitably need to supplant the ecosystem reserves once the initial AAVE allocation in the ecosystem reserves is exhausted (either by the DAO or by reward claims). While the Aave risk docs suggest that the reserve factor is primarily calibrated against the "riskiness" of a token, we propose that a crucial factor for the reserve factor is understanding the elasticity of borrowers and suppliers.

We suggest optimizing protocol reserves benefits both protocol risk management and the DAO's business model. Increasing the reserve factor excessively could harm the protocol. As reserve factors increase, elastic suppliers exit the system, leading to high utilization risk and a decrease in overall revenue if borrowers also leave. However, beyond a certain risk threshold - slightly below the optimal utilization - decreasing utilization is strictly bad for the protocol: the protocol pays more suppliers the cost of capital for a token that isn't being lent out.

The protocol reserves currently comprise DAI, USDT, and WETH accrual (on Ethereum). These, alongside WBTC and USDC, will be the focal assets in the recommendations presented in this paper.

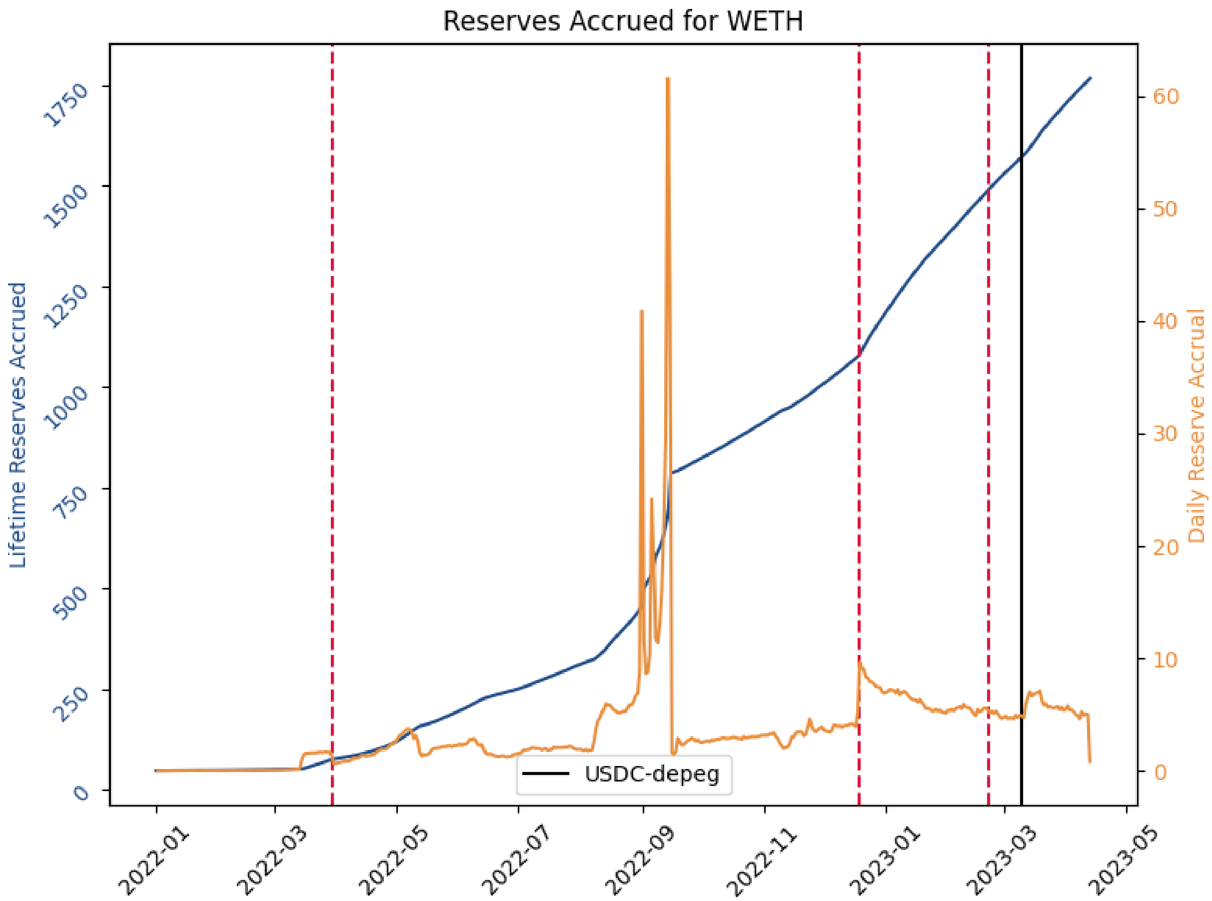

Figure 3.4: Reserve accrual is likely sensitive to parameter changes. Raising the reserve factor may result in an immediate increase in accruals, but in the long term may lead to a decrease in supply and, therefore, an increase in borrow rates. This borrow rate increase may lead to a general decrease in borrowing activity, returning reserve accrual to its original levels or below it. This is likely what we observed with WETH in December 2023.

2.2 Elasticity

Elasticity is a key concept when examining changes in supply and demand within the Aave protocol. In economics, elasticity represents the degree to which the quantity demanded or supplied responds to variations in price or interest rates. A high elasticity indicates a significant change in quantity demanded or supplied due to a price or interest rate change, whereas a low elasticity indicates a less significant change.

In the context of Aave, we are interested in understanding how elastic the borrowers and suppliers are to alterations in interest rates. If the borrowers and suppliers exhibit high elasticity, minor adjustments in interest rates could trigger considerable changes in borrowing and supply behavior. Conversely, if inelastic, their actions would change less in response to the same rate changes.

Here's how we might assess supplier and borrower elasticity under different scenarios:

Scenario 1: A shift in supply results in a change in the supply rate. If suppliers are elastic, they will modify their behavior to counterbalance this change and re-establish the supply rate at equilibrium. For instance, if a large withdrawal occurs, it may be succeeded by a substantial deposit. To some extent, we can measure this type of elasticity by performing mean reversion tests.

Scenario 2: A shift in borrowing leads to a change in the supply rate. If suppliers are elastic, they will adjust their behavior again to counteract this discrepancy and return the rates to equilibrium. We can evaluate this form of elasticity using correlation tests, cross-correlation functions (which account for lagged effects), or Granger causality tests. We can express this by denoting D as the quantity borrowed, $Q(D)$ as the quantity supplied, and ϵ_S as a candidate measure of supplier elasticity.

$$\epsilon_S = \frac{\partial Q(D)}{\partial D} \frac{D}{Q}$$

The aforementioned methodologies can also be employed to gauge the elasticity of borrowers. However, irregular supply and demand data can obfuscate such data analysis, particularly in low-utilization markets with fewer borrowers and suppliers. We have contemplated and employed various tests for the relationships mentioned above in our investigation. These will aid us in interpreting supply and demand graphs and hint at which potential outcomes in our parameter recommendations are more probable. In many instances, statistical tests do not instill confidence in their findings. It's vital to remember that the absence of evidence is not evidence of absence. As we've discussed, a useful mental model is that relative elasticity explains who is more steadfast in their current equilibrium rate. Thus, the relative elasticity of borrowers and suppliers can provide insights into who is more likely to persist in their current borrowing or supplying behavior despite interest rate changes. If borrowers exhibit higher elasticity, they are more likely to adhere to the current borrowing rate than suppliers are to stick to the current supply rate. Suppose the utilization is high, and we intend to reduce it by lowering the reserve factor. In that case, the current supply and borrow rates will be disturbed, leading to whoever is more elastic modifying their behavior and nudging utilization towards their equilibrium rate.

If borrowers are more elastic, they will adhere to the current borrowing rate, which happens at the exact utilization as before the reserve factor adjustment, despite the elevated supply APR. This suggests we might not witness a substantial decrease in utilization because borrowers stick to the current borrowing rate. In such a case, reducing the reserve factor

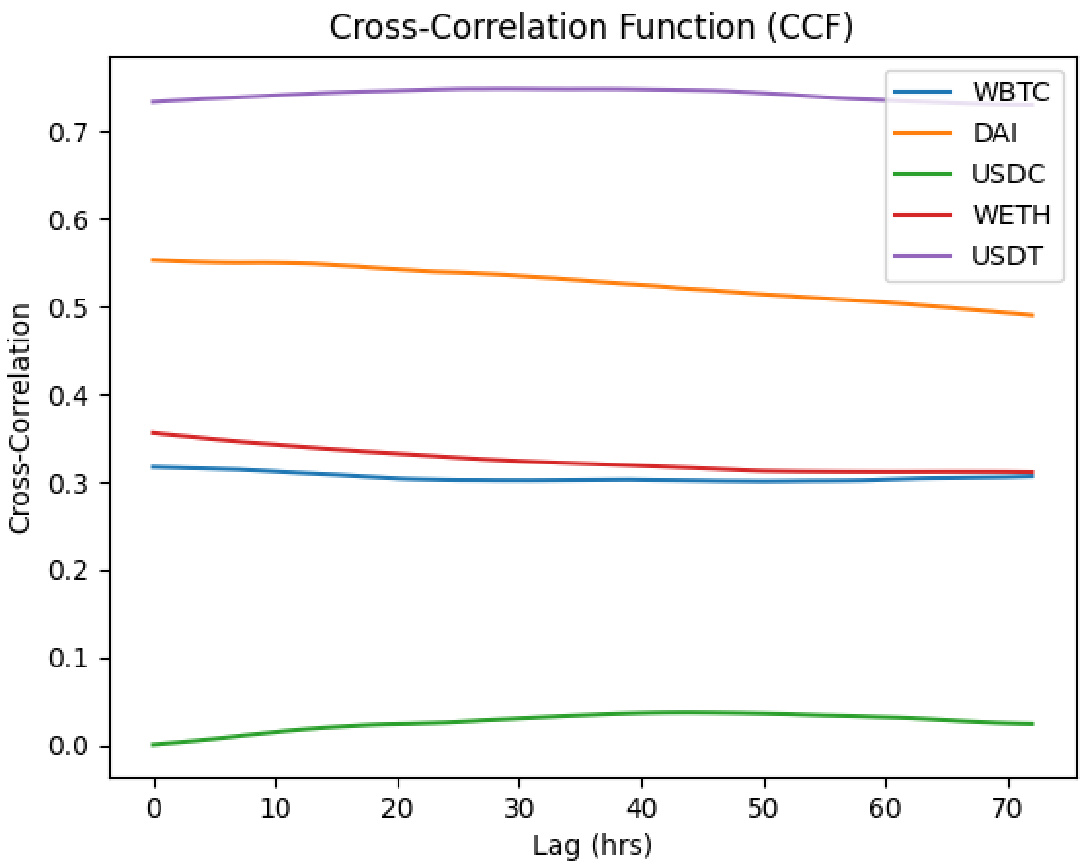

Figure 3.5: Cross-correlation testing on various Aave v2 markets using data from October 2022 to April 2023. Cross-correlation functions suggest that USDT borrowers are more elastic to quantity supplied (and vice versa) than other tokens. While these graphs may help build some intuition on how borrower and supplier elasticity varies by token, token supply and demand is non-stationary so CCF results may be spurious. Accordingly, we do not leverage these results explicitly in our framework.

might negatively impact the protocol since the utilization wasn't significantly affected, but the protocol's reserve accrual was decreased.

These concepts and their implications will be further examined in the results section.

2.3 Granger Causality Testing

Granger causality tests assist us in establishing whether an independent time series (such as utilization) can forecast a dependent time series (like the borrowed amount). Essentially, if alterations in utilization Granger cause changes in the borrowed amount, it indicates that borrowers are predictably responsive to variations in utilization. For instance, a drop in utilization (potentially due to an increase in supply) may predictably trigger a rise in borrows. This can also be applied to supply dynamics.

Although not a perfect measure of elasticity, Granger causality provides statistical evidence to back claims about whether borrowers or suppliers are more responsive to changes in utilization (and, therefore, interest rates).

Assuming borrowers are elastic (and more so than suppliers), we could reduce the first slope of the interest rate curve and anticipate a net increase in borrows. This rise in borrows could offset the decrease in the borrowing rate, possibly leading to a net increase in protocol revenue. However, if reducing the slope only leads to a net outflow of supply, borrow rates may remain below their current equilibrium rate without provoking a substantial increase in net borrows, resulting in a revenue loss.

For accurate application of Granger causality testing, we need to assume the following:

- **Stationarity of time series:** Instead of testing borrow and supply amounts directly, we test the differences (i.e., changes) in these amounts. We must verify that these differences are stationary using the Augmented Dickey-Fuller test.
- **Linear response:** We assume that alterations in interest rates result in linear changes in borrow and supply amounts.
- **Absence of structural breaks:** Structural breaks refer to sudden shifts in borrow and supply activity distribution. We contend that AIPs that modify interest rate parameters are the only true structural breaks we observe.

2.4 Interest Rate Parameters and Utilization

Given these considerations, we can propose a framework for setting interest rate parameters and the reserve factor to optimize optimal utilization. This approach will vary depending on whether a pool is over-utilized, under-utilized, or near-optimal utilization.

The Utilization Playbook

We propose incremental changes that allow us to gather data and mitigate the risk that our elasticity measurements were incorrect or that exogenous changes to market rates unexpected changes to supply and/or demand. Given that utilization may be inherently volatile, especially in smaller markets, we do not target optimal utilization itself. Instead, we create a risk buffer to avoid exceeding optimal utilization by targeting one standard deviation below optimal. Let σ_u be the standard deviation in utilization in our lookback period, then:

$$u_T = u_{\text{opt}} - \sigma_u$$

is our target utilization.

When it comes to utilization, there are generally three scenarios to consider:

1.Current utilization is at the target: - This is the ideal state where capital efficiency is maximized, a lower spread between borrowers and suppliers, and a higher revenue share goes to the protocol. In such a scenario, it may be appropriate to raise the base interest rate to reduce interest rate volatility or increase the second slope parameter if we observe prolonged periods of high utilization. 2. **Current utilization is ¿ 1 standard deviation below the target:** Under-utilization is unsuitable for both the protocol and the suppliers as the interest paid by borrowers gets diluted among many suppliers despite most of the supply being left unused. In such cases, an increase in the reserve factor, or decrease in r_1 and r_2 might be appropriate to encourage more borrowing or disincentivize supply. 3. **Current utilization above the target:** Over-utilization can pose a risk to the protocol, increasing the likelihood of missed liquidations and the chance of suppliers' inability to withdraw their funds. For sustained periods of over-utilization, it might be beneficial to increase in r_1 and decreases in the reserve factor, as well as potential increases to r_2.

Elasticity is a valuable tool that informs us about the most probable outcome once changes have been implemented. In other words, it can help predict whether we will likely see utilization trending towards the supplier's equilibrium rate, the borrower's equilibrium rate, or if utilization will remain unchanged.

Targeting Borrow and Supply Rates

When it comes to targeting borrow and supply rates, it is essential to note that we target borrow rates with changes to the first slope parameter instead of supply rates. Lowering the slope to match the supply rate is almost certainly a bad idea as it would result in lower borrowing rates at optimal utilization than the current borrowing rate. This would lead to an unstable equilibrium around the optimal utilization, and if borrowers are elastic to the current equilibrium borrow rate, they might consider borrowing beyond the optimal level. Recall our borrow and supply rates at equilibrium u^*:

$$r_b^* = r_0 + \frac{u^*}{u_{\text{opt}}} r_1$$

$$r_s^* = u^* \left(r_0 + \frac{u^*}{u_{\text{opt}}} r_1 \right) (1 - RF)$$

If we lower $r_1 \to r_1'$ such that $r_s'(u_{\text{opt}}) = r_s^*$, we'd find:

$$r_1' = \frac{r_s^*}{u_{\text{opt}}(1 - RF)} - r_0$$

In the best case, we achieve optimal utilization, meaning we get that the borrowing rate is at most:

$$r_b' = \frac{r_s^*}{u_{\text{opt}}(1 - RF)} \tag{3.7}$$

Substitute Equation 12 in Equation 14 we find:

$$r_b' = \frac{u^* \left(r_b^*\right)}{u_{\text{opt}}}$$

Since $u^* < u_{\text{opt}}$, we find that the new borrow rate at optimal utilization would still be lower than the current equilibrium borrow rate. This change would require a significant increase in borrowing activity to prevent a material decrease in protocol revenue, implying that we need significant borrowing elasticity.

Instead, we argue that the slope should be lowered to match the current equilibrium borrow rate. If suppliers are more elastic, we target them with the reserve factor:

$$r_s(u_T)$$

2.5 Reserve Factor as a Tax

The reserve factor can be understood as a form of tax applied at the supplier level. Both borrowers and suppliers may bear the impact of this tax: if suppliers earn a lower APR, they may supply less, which leads to higher utilization and, consequently, a higher borrowing rate. From basic economic theory, we know that the tax incidence, which is essentially who bears the most tax burden, is inversely proportional to elasticity. In other words, if suppliers are more elastic than borrowers, then borrowers will bear most of the tax burden, and vice versa.

Understanding the tax incidence can help predict how a change in the reserve factor will impact utilization. For instance, if the utilization is currently low and we increase the reserve factor, and if suppliers are inelastic, they would bear the full burden of the tax, resulting in a lower interest rate for them, with the difference being redirected to the protocol. In this scenario, utilization remains unchanged. While this could benefit the protocol's revenue, it wouldn't help our goal to increase utilization.

If suppliers are elastic and borrowers are inelastic, borrowers would bear the burden of the increased reserve factor, resulting in a higher borrowing rate. In this case, utilization increases to maintain the equilibrium supply rate, which now occurs at a higher utilization level.

If both borrowers and suppliers are elastic, the new utilization rate would be somewhere between the current utilization (where the equilibrium borrowing rate still occurs) and the utilization where the old equilibrium supply rate occurs. The extent of the increase in utilization would then depend on the relative elasticity of supply and demand.

2.6 Under-Utilization

Under-utilization often signifies an over-supply scenario: a significant portion of the protocol's revenue is spread thinly amongst too many suppliers, even though most of their supply remains unused. To address this and aim for higher utilization, various strategies can be employed depending on certain conditions, such as the elasticity of borrowers and suppliers. **1. Lower the base rate r_0, if $r_0 > 0$:** This strategy steeps the interest rate curve, causing the current equilibrium rates for borrow and supply to occur at higher utilizations. If either borrowers or suppliers are elastic to their respective rates, we can anticipate an increase in utilization, subsequently increasing the protocol's reserve accrual. However, if neither borrowers nor suppliers are elastic, we may not see an increase in utilization, leading to lower borrowing, supply, and protocol rates.

Condition: Either borrowers or suppliers are elastic.

2. Lower the slope r_1: This action ensures that the current equilibrium borrow rate occurs slightly below optimal, instantly reducing the borrowing rate at the current utilization. If borrowers are elastic, we might see an increase in borrowing, bringing the borrowing rate back to equilibrium at higher utilization. If borrowers are inelastic, we may not see an increased borrowed amount. If suppliers are inelastic, the amount supplied might decrease, leading to higher utilization. However, this could still result in a lower APR for the protocol.

Condition: Borrowers are elastic.

3. Raise the reserve factor: This action ensures that the supply equilibrium rate occurs slightly below optimal, instantly reducing the supply rate at the current utilization. If suppliers are elastic, we might see a decrease in supply, leading to increased utilization and, subsequently, higher borrowing and supply rates. If borrowers are inelastic, we achieve our goal of increasing utilization with borrowers paying a higher rate. However, if borrowers are also elastic, they might reduce utilization to match their old equilibrium borrowing rate.

Condition: Either borrowers or suppliers are inelastic.

If suppliers are inelastic, utilization does not increase, but the protocol increases its APR. If suppliers are elastic, but borrowers are inelastic, then utilization increases, and the borrowers pay a higher rate to the protocol.

In most scenarios, a combination of strategies (2) and (3) would be suggested to balance the risks associated with inaccurate estimates of borrower and supplier elasticity. Moreover, considering that equilibrium borrow and supply rates fluctuate over time, it is advisable to favor moderate changes to avoid risks associated with equilibrium rates rising past optimal utilization.

2.7 Over-Utilization

Over-utilization signifies a scenario where equilibrium occurs at or above the optimal utilization level, indicating that market borrow and supply rates are higher than the current optimal. Alternatively, over-utilization might manifest with the pool's utilization below optimal but with frequent, sustained periods exceeding the optimal due to frequent and large liquidations or sporadic withdrawals by a few major suppliers. These scenarios require separate considerations and strategies.

In case of consistent over-utilization above optimal, strategies like lowering the reserve factor or raising r_1 can be applied:

1. Raising the slope r_1: This strategy increases borrowing costs such that the current borrow rate occurs below optimal utilization. If borrowers are elastic to the current borrowing rate, they will reduce their borrowing activity, bringing utilization below the optimal level. While this method increases the net interest charged per borrow, it decreases the net amount being borrowed.

2. Lowering the RF (reserve factor): This strategy increases the supplier APR such that the current supply rate equilibrium occurs below optimal utilization. If suppliers are elastic, this leads to a lower utilization without necessarily changing the net amount being borrowed. However, if suppliers are more elastic than borrowers, there might be an increase in both borrows and supply without a significant change in utilization.

Condition: Suppliers are more elastic than borrowers.

In such scenarios, it's recommended to raise r_1, regardless of borrower or supplier elasticity. If borrowers and suppliers are inelastic, the protocol will still observe a net increase in reserve accrual. If either borrowers or suppliers are elastic, there will be a decrease in utilization. When periods of over-utilization occur due to supplier withdrawals, it could be that the equilibrium borrow and supply rates are below optimal utilization. Still, there's insufficient elasticity in both borrowing and supply to correct utilization after a large withdrawal. Essentially, borrowers and suppliers aren't responding quickly enough to increases in utilization. In such cases, it is suggested to raise r_2 to increase incentives for deposits or repayments.

2.8 Advocating for Incremental Adjustments

Regardless of the specific scenario in which modifications may be deemed necessary, our strategy is to aim for either the equilibrium borrow rate or the equilibrium supply rate to materialize slightly beneath optimal utilization. Recognizing that utilization can be volatile, especially in niche markets, we advocate for a policy of progressive, moderate alterations to interest rate parameters. Such an incremental approach enables us to accumulate valuable data pertaining to user responsiveness, technically referred to as elasticity, in relation to shifts in parameters. Moderate adjustments are recommended to mitigate the potential risk of utilization sporadically peaking beyond the optimal range, especially if prevailing market borrow and supply rates demonstrate an upward trajectory.

Let σ_u denote a standard deviation in utilization.

Example - Lowering r_1: We target the equilibrium borrow rate r_b^* to occur one standard deviation in utilization below optimal. That is, $r_b(u_{\text{opt}} - \sigma_u) = r_b^*$. However, as market rates and utilization volatility may change quickly, we de-risk our changes by gradually targeting optimal utilization. For example, WBTC optimal utilization is 65% with recent volatility of 2% and borrow rates fluctuating closely around 1.2%. We would then target $r_b(0.63) = 0.012$, resulting in a slope $r_1 \approx 0.0124$. Implementing a substantial reduction in the slope carries inherent risks. If borrowers demonstrate inelasticity towards borrow rates (for instance, if there is no supplementary demand for borrowing bitcoin beyond the existing borrowers), we could potentially face a significant diminishment in the protocol's reserve accrual. Instead, we propose gradually altering WBTC's slope from 8% to 6%. This approach enables us to accumulate empirical data on borrower elasticity. We propose an increase in the reserve factor to mitigate potential detriments to the protocol. As the subsequent analysis will demonstrate, WBTC is an over-supplied market and could benefit from a significant reduction in supply.

2.9 Utilization Trends

Our analysis considers utilization and interest rate data within three to six-month intervals. Taking into account more extended periods could introduce bias against recent rate trends. As a potential enhancement to our methodology, we might consider giving more weight to recent observations, possibly through exponentially weighted moving averages or other techniques that favor recent data.

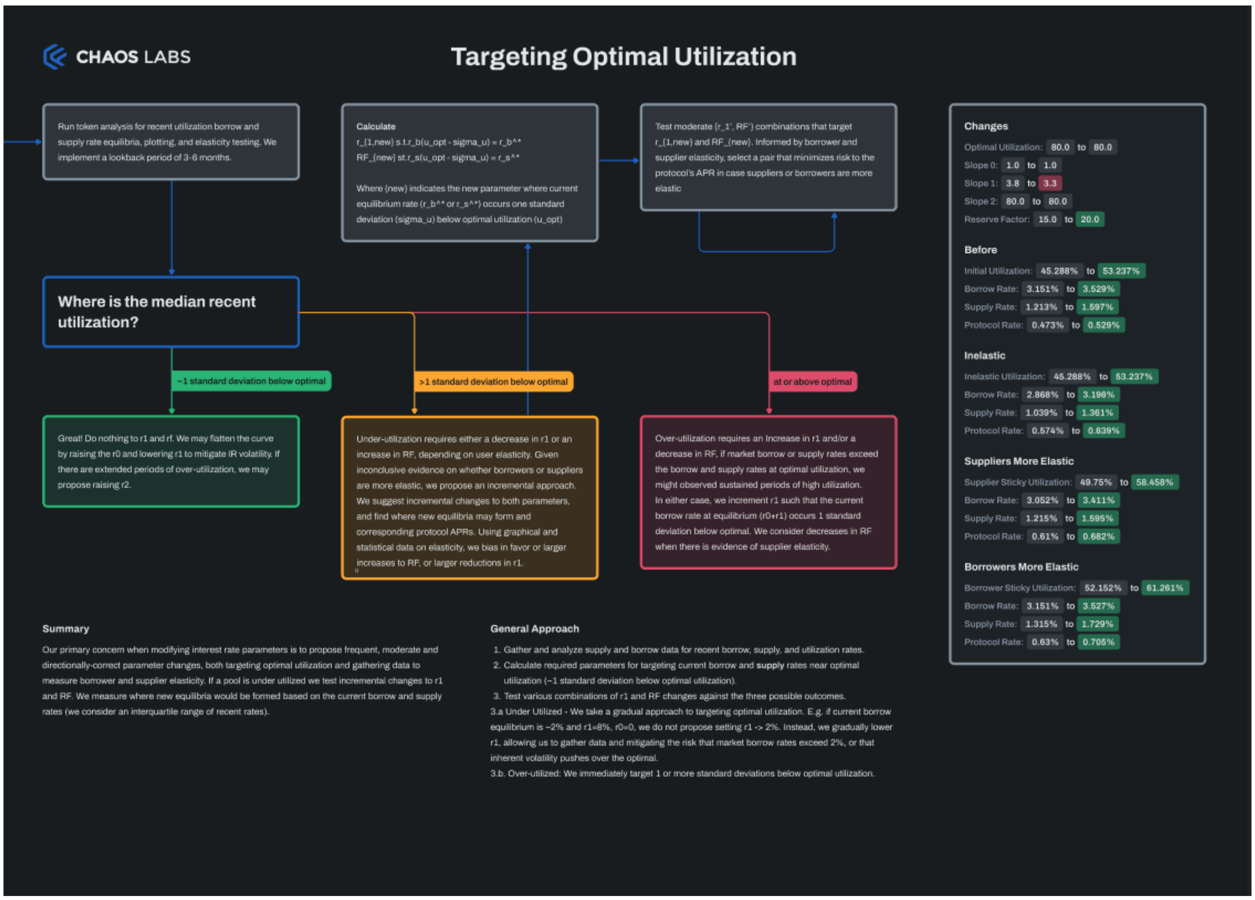

Figure 3.6: Flow chart which summarizes the recommended methodology.

Chapter 4

Results

1 Results

Market rates are subject to rapid fluctuations, and developing a framework to predict the "competitive equilibrium market rates" for each token/market/deployment presents a complex challenge that may overtax risk providers, strain governance, and, in all likelihood, yield inaccurate predictions. We advocate for an approach that observes the rates borrowers and suppliers are willing to accept on Aave and aims for optimal utilization in accordance with these rates. Implementing an iterative approach characterized by frequent, transparent, and straightforward modifications to interest rates will enable the Aave community to understand the elasticity of borrowers and suppliers better while targeting optimal utilization.

In formulating our recommendations, we strive to mitigate the risk of substantial market changes in borrowing and supply rates in the weeks after implementing a parameter change. For instance, the emergence of a new yield farming opportunity for a particular token could instigate a significant surge in the market borrowing rate. If, instead of adopting an iterative approach to targeting optimal utilization, we hastily compel the current market borrow rate to match optimal utilization, this abrupt increase in market borrow rate could thrust utilization beyond the optimal point and expose the protocol to considerable risk.

We propose the following utilization target:

$$u_T = u_{\text{opt}} - \sigma_u$$

where σ_u is the recently observed standard deviation in utilization. We target one standard below optimal to avoid exceeding optimal utilization due to inherent volatility.

1.1 WBTC on Ethereum v2

This section applies our methodology to Wrapped Bitcoin (WBTC) on the Aave v2 Ethereum market. We assess WBTC for elasticity from December 2022 to April 2023. Following the steps outlined in the previous flowchart, we first establish optimal utilization based on collateral at risk and then target this optimal utilization by adjusting the parameters r_{slope1} and RF, considering the elasticities of supply and demand.

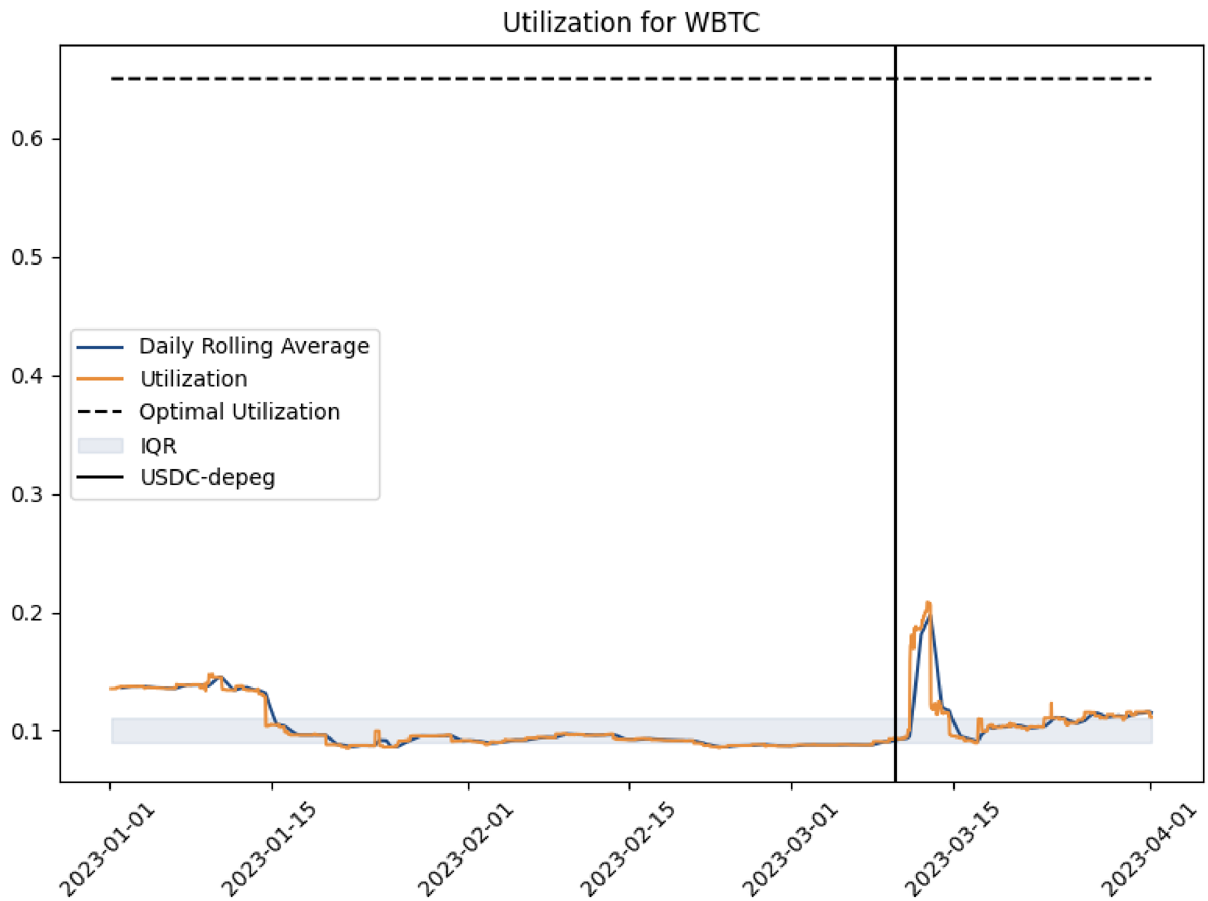

Figure 4.1: WBTC utilization, notice that it is significantly below optimal, meaning there is an over-supply of WBTC diluting the protocol's revenue from this market.

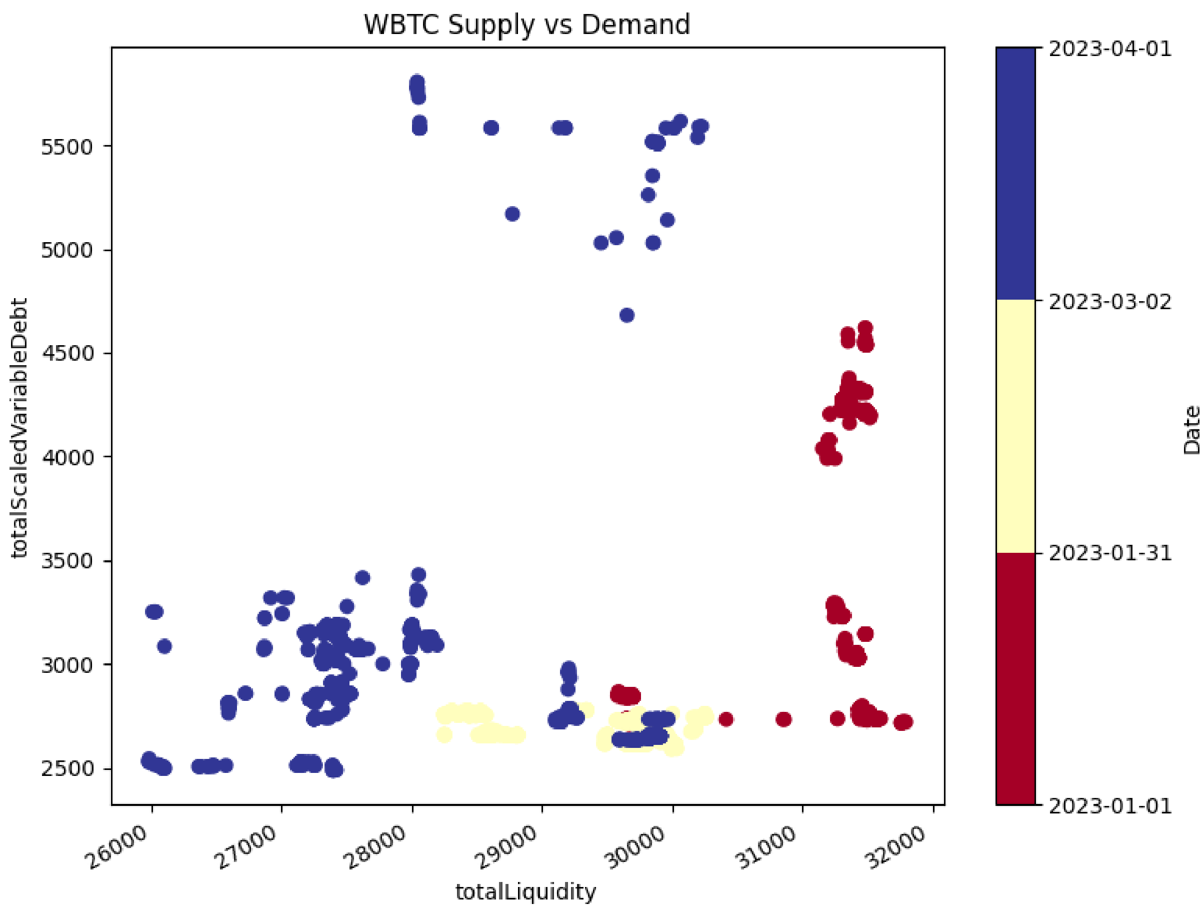

Figure 4.2: A plot of variable debt and protocol liquidity. Positive sloping diagonal lines indicate elasticity between borrowers and suppliers (e.g., an increase in supply leads to lower borrowing rates which leads to an increase in borrows). Vertical lines (e.g., January 2023) and horizontal lines (e.g., Feb 2023) indicate supplier or borrower inelasticity.

u_{opt} : $0.65 \rightarrow 0.75$ (0.77 rounded down), based on collateral at risk. This flattens the curve - encouraging borrowing while decreasing current rates - and improves potential capital efficiency for WBTC. Furthermore, the pool is significantly under-utilized, and we do not observe significant elasticity in supply and demand:

symbol	start	end	supply is elastic	supply elasticity stat	lag	demand is elastic	demand elasticity stat	lag
WBTC	2023-01-01 00:00:11	2023-01-31 00:00:11	False	0.036	3	False	1.013	3
WBTC	2023-01-31 00:00:11	2023-03-02 00:00:11	False	1.739	1	False	0.008	1
WBTC	2023-03-02 00:00:11	2023-04-01 00:00:11	True	7.580	22	True	2.132	22

Figure 4.3: Granger-causality testing output. We cannot find statistically significant evidence of intra-month elasticity in WBTC outside of March (which is particularly noisy due to the depeg in USDC). Notice that this is visually corroborated by the supply and demand chart above.

We find the 3-month median borrow rate of 1.2%

```
============== CHANGES =================

Optimal Utilization: 65.0 -> 65.0
Slope 0: 0.0 -> 0.0
Slope 1: 8.0 -> 6.0
Slope 2: 300.0 -> 300.0
Reserve Factor: 20.0 -> 30.0

=============== BEFORE =================

Initial Utilization: 9.017% to 11.084%
Borrow Rate: 1.11% to 1.364%
Supply Rate: 0.08% to 0.121%
Protocol Rate: 0.222% to 0.273%

============== INELASTIC ===============

Inelastic Utilization: 9.017% to 11.084%
Borrow Rate: 0.832% to 1.023%
Supply Rate: 0.053% to 0.079%
Protocol Rate: 0.25% to 0.307%

====== SUPPLIERS MORE ELASTIC =========

Supplier sticky utilization: 11.111% to 13.714%
Borrow Rate: 1.026% to 1.266%
Supply Rate: 0.08% to 0.122%
Protocol Rate: 0.308% to 0.38%

====== BORROWERS MORE ELASTIC =========

Borrower sticky utilization: 12.012% to 14.815%
Borrow Rate: 1.109% to 1.368%
Supply Rate: 0.093% to 0.142%
Protocol Rate: 0.333% to 0.41%
```


Figure 4.4

Despite their magnitude (i.e., an increase in the reserve factor from 20% to 30%), these adjustments are unlikely to pose substantial risks to the protocol, considering that the current equilibrium borrowing and supplying rates can be achieved with comparatively minor shifts in utilization. This preliminary proposal aims to evaluate this hypothesis, providing a basis for proposing more consequential modifications once relevant data has been collected, particularly for large but inelastic markets. In general, the likelihood of experiencing a detrimental impact on reserves is low, while the potential for a minor increase in utilization and possibly borrowing activity is relatively high.

In situations where either borrowers or suppliers demonstrate elasticity, as evidenced toward the end of Q1 2023, we anticipate an uptick in utilization and a resultant enhancement in the pool's capital efficiency. Given the minimal volatility in utilization, there is no compelling reason to increase the base slope. Similarly, no evidence suggests that the current slope 2 exerts any influence on WBTC utilization, so we do not recommend any modifications in this regard.

1.2 WBTC on Ethereum V3

WBTC utilization exhibits a similar trend of being low, although it is worth noting that WBTC users appear to demonstrate significantly higher elasticity as more borrowers and suppliers transition to v3. Consequently, we propose aligning the WBTC parameters on v2 and v3 to target higher utilization levels across both platforms.

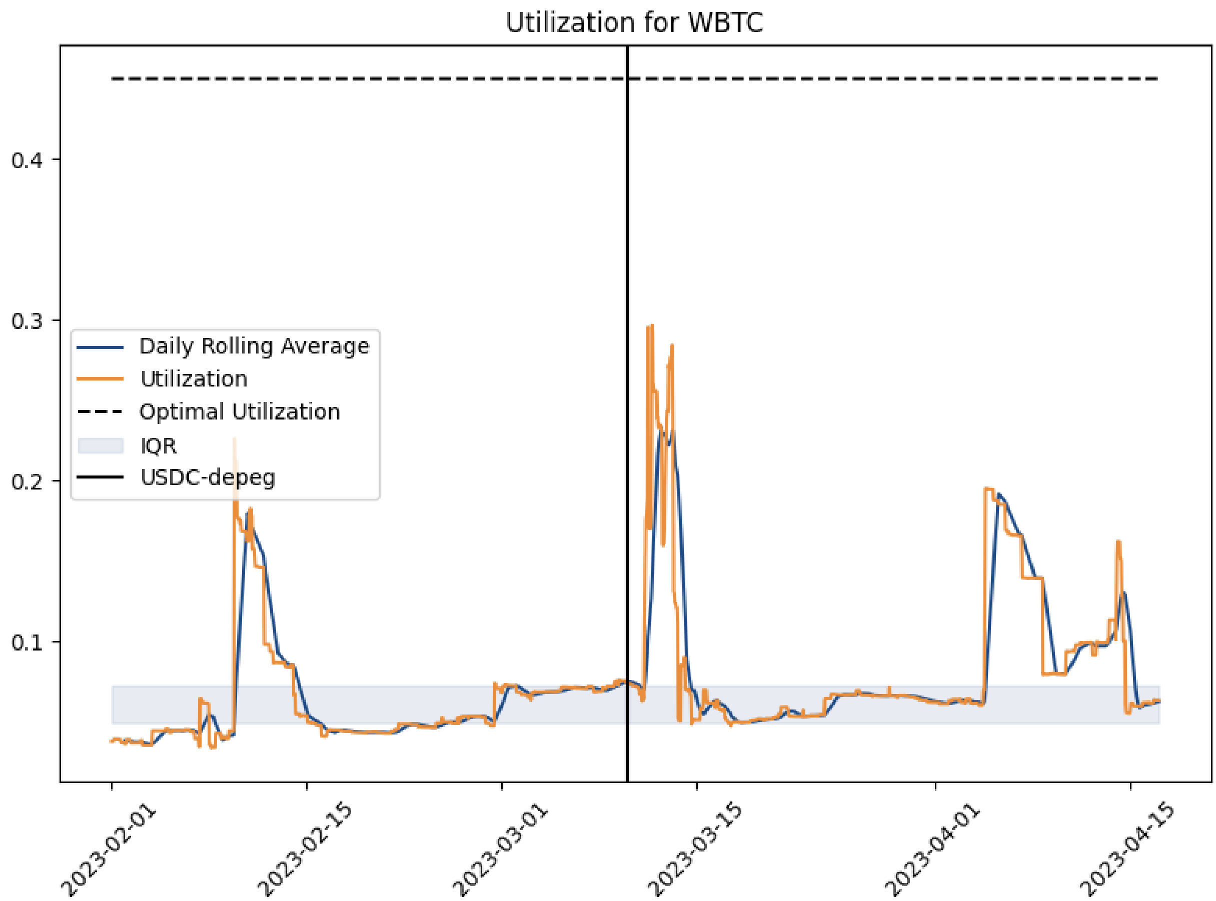
Utilization for WBTC
0.4
0.3
0.2
0.1
Daily Rolling Average
Utilization
Optimal Utilization
IQR
USDC-depeg
2023-02-01
2023-02-15
2023-03-01
2023-03-15
2023-04-01
2023-04-15

Figure 4.5

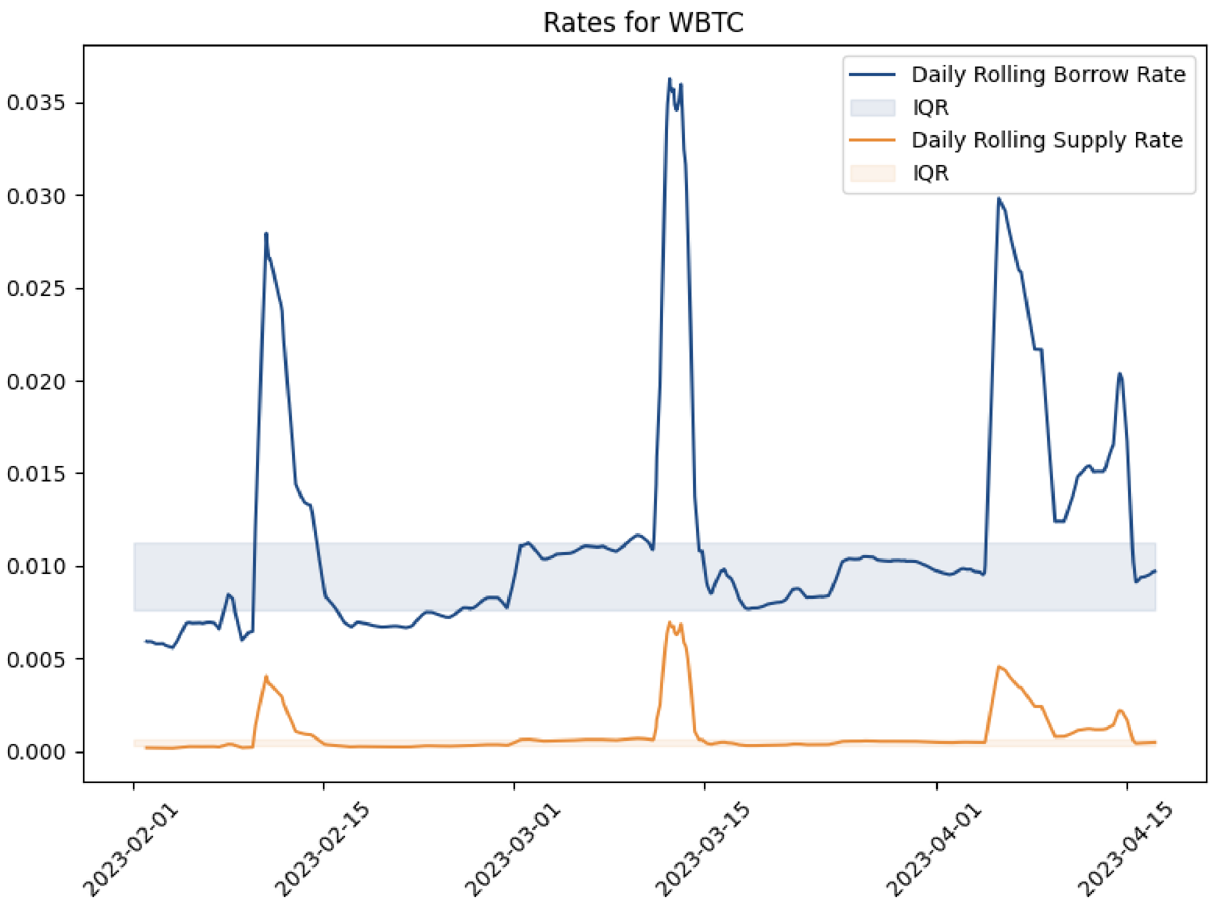
Rates for WBTC
Daily Rolling Borrow Rate
IQR
Daily Rolling Supply Rate
IQR
0.035
0.030
0.025
0.020
0.015
0.010
0.005
0.000
2023-02-01
2023-02-15
2023-03-01
2023-03-15
2023-04-01
2023-04-15

Figure 4.6

1.3 WETH on Ethereum v2

The WETH market on Aave v2 has witnessed a substantial increase in both borrower and supplier elasticity throughout 2022 and early 2023, as demonstrated both graphically and via our statistical analysis. Despite the relatively higher utilization rate of WETH as compared to WBTC, it consistently remains 20-25% below the optimal rate, fluctuating between 45% and 53% from January to April 2023.

As discussed in the introductory section, the previous proposal entailed both an increase in the base rate (from 0% to 1%) and a reduction in r_{slope1}. This resulted in mixed effects on the borrowing rates (and consequently the supply rates): borrowing costs increase when utilization is below 46%, but decrease when utilization exceeds 46%. Given the observed utilization range between 45% and 53%, no significant change in the borrowing rate has been recorded. Consequently, we have not seen substantial improvement in either utilization or reserve accrual.

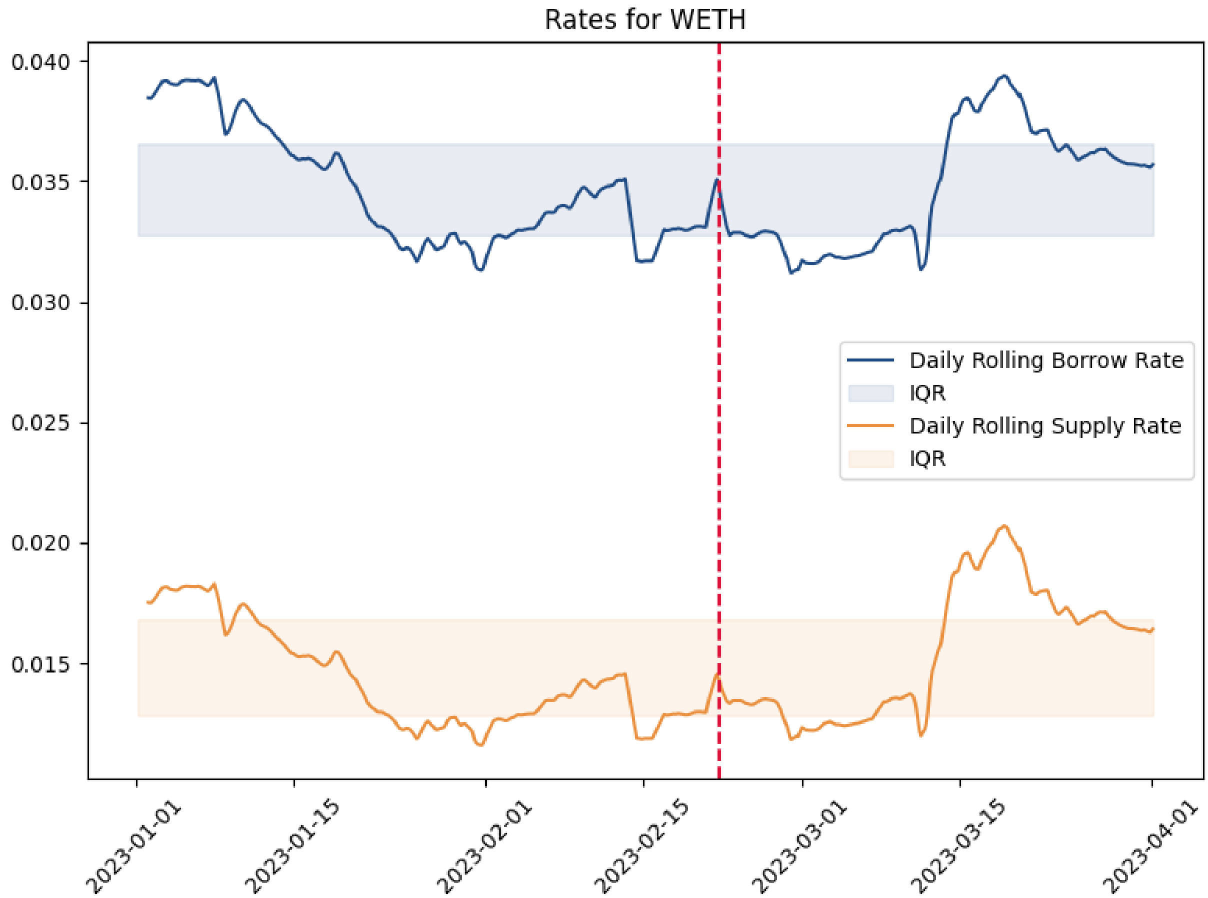

Figure 4.7: The red dotted line indicates a change in interest rate parameters. Recall that the USDC depeg occurred around March 10th, 2023.

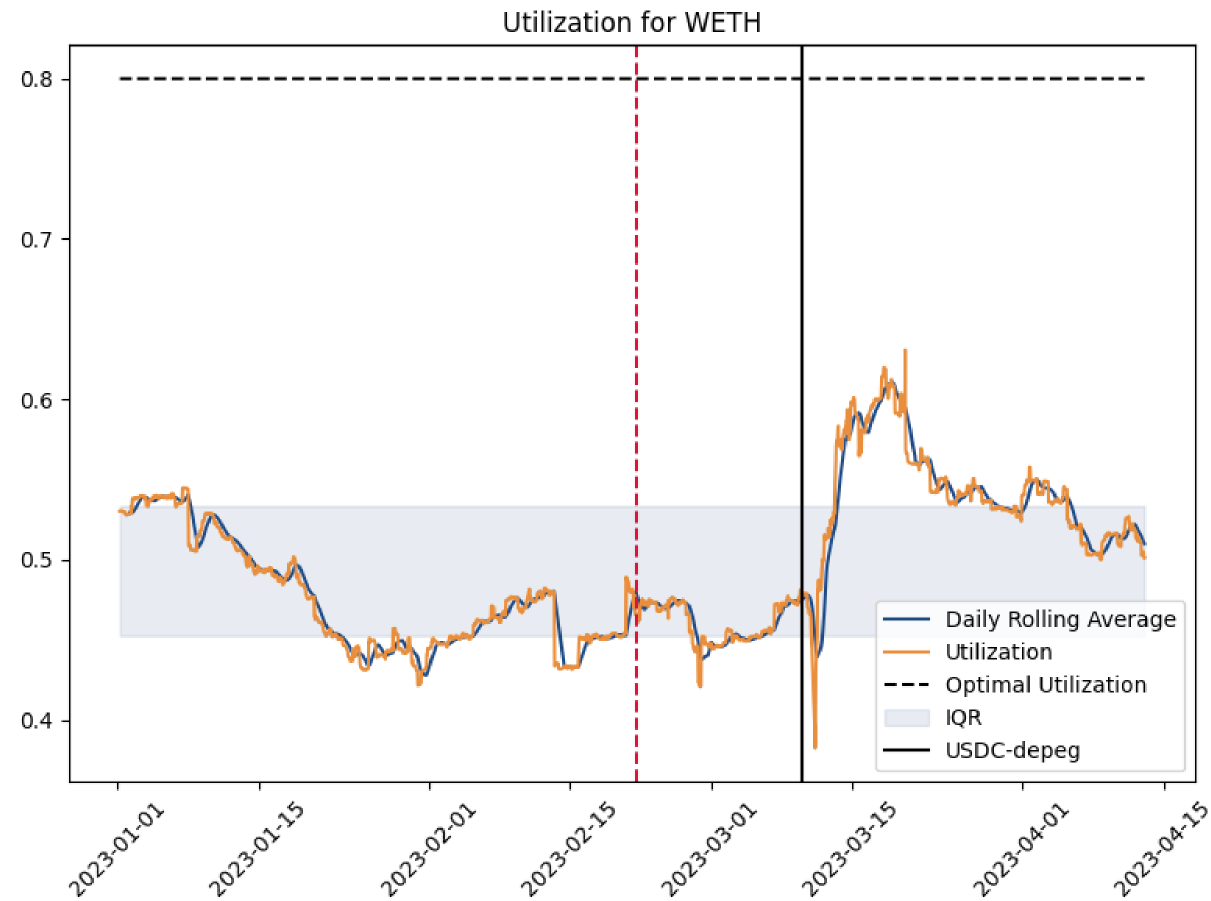

Figure 4.8: WETH on Aave v2 Ethereum utilization.

In light of the fact that WETH is an appealing borrowing choice (among the top three revenue generators for Aave), it is plausible that a significant reduction in borrowing rates at the current utilization level could gradually drive up utilization over time. To encourage greater borrowing and increase utilization, we propose a two-part approach. Firstly, we suggest lowering the base rate, which effectively steepens the interest rate curve, pushing the current equilibrium rates toward the optimal rate. Additionally, considering the oversupply in the WETH market and the likely greater elasticity of borrowers compared to suppliers, there is an opportunity to contemplate raising the reserve factor:

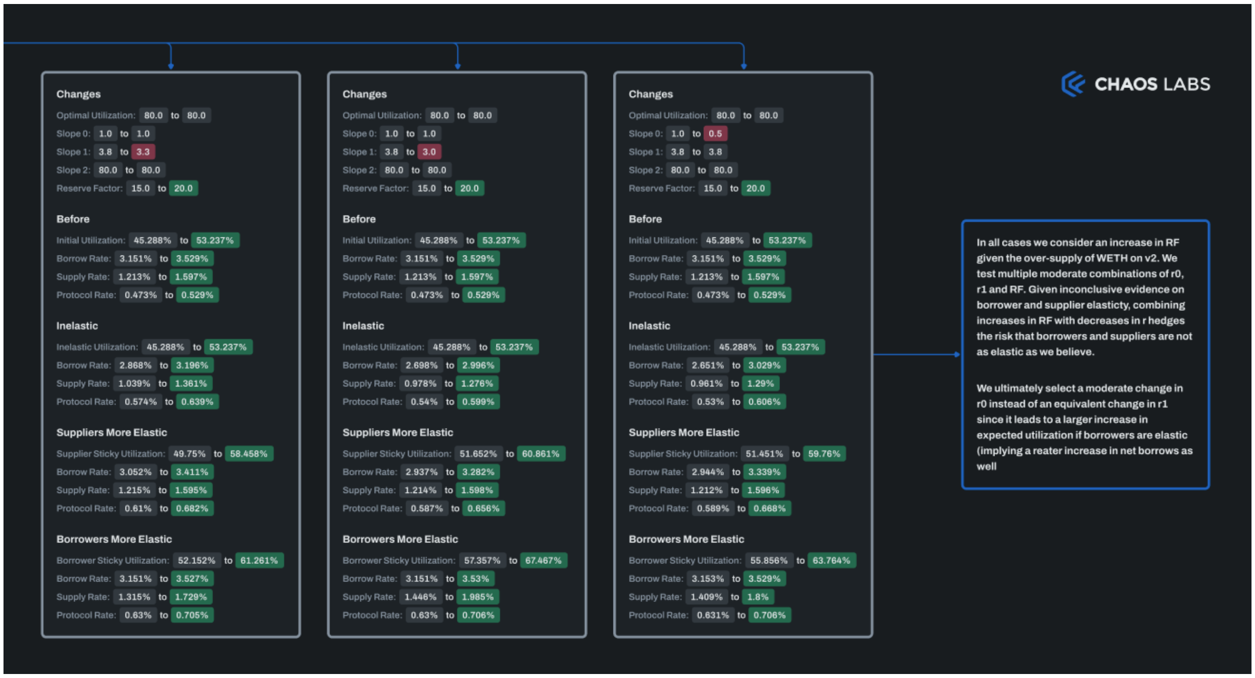

Figure 4.9

These proposed modifications should bring about moderate increases in utilization while also managing the risk of unexpected surges past the optimal rate. As discussed, such a scenario could arise due to exogenous shocks impacting market borrowing or supply rates. Implementing a strategy that combines a reduction in the base rate with a moderate increase in the reserve factor could result in significant enhancements in the protocol's reserves. If borrowers are more elastic, while the "Protocol Rate" remains unchanged (since the borrowing rate remains flat), the protocol's net reserve accrual would increase. In other words, if the rate curve is lowered, borrower elasticity implies an increase in net borrows. While the protocol's rate remains constant, the actual revenue will see an upturn:

$$\text{Protocol Reserve Accrual} = B \cdot r_b \cdot RF \tag{4.1}$$

2 WETH on Ethereum V3

WETH on Aave v3's Ethereum platform exhibits more elastic behavior for borrowers and suppliers, with gross borrows and deposits generally increasing since the pool's inception in early 2023.

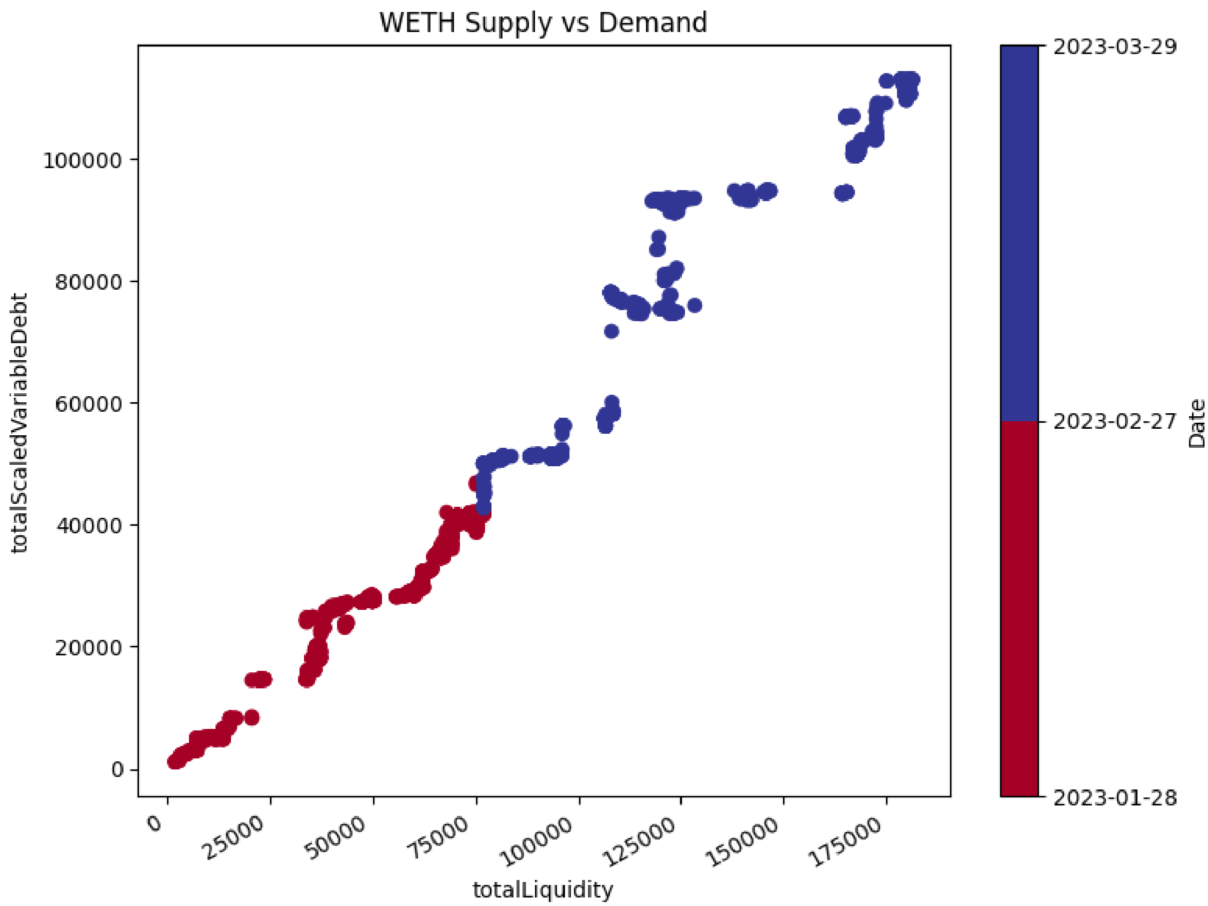
WETH Supply vs Demand
totalScaledVariableDebt
totalLiquidity
0
20000
40000
60000
80000
100000
0
25000
50000
75000
100000
125000
150000
175000
2023-03-29
2023-02-27
2023-01-28
Date

Figure 4.10

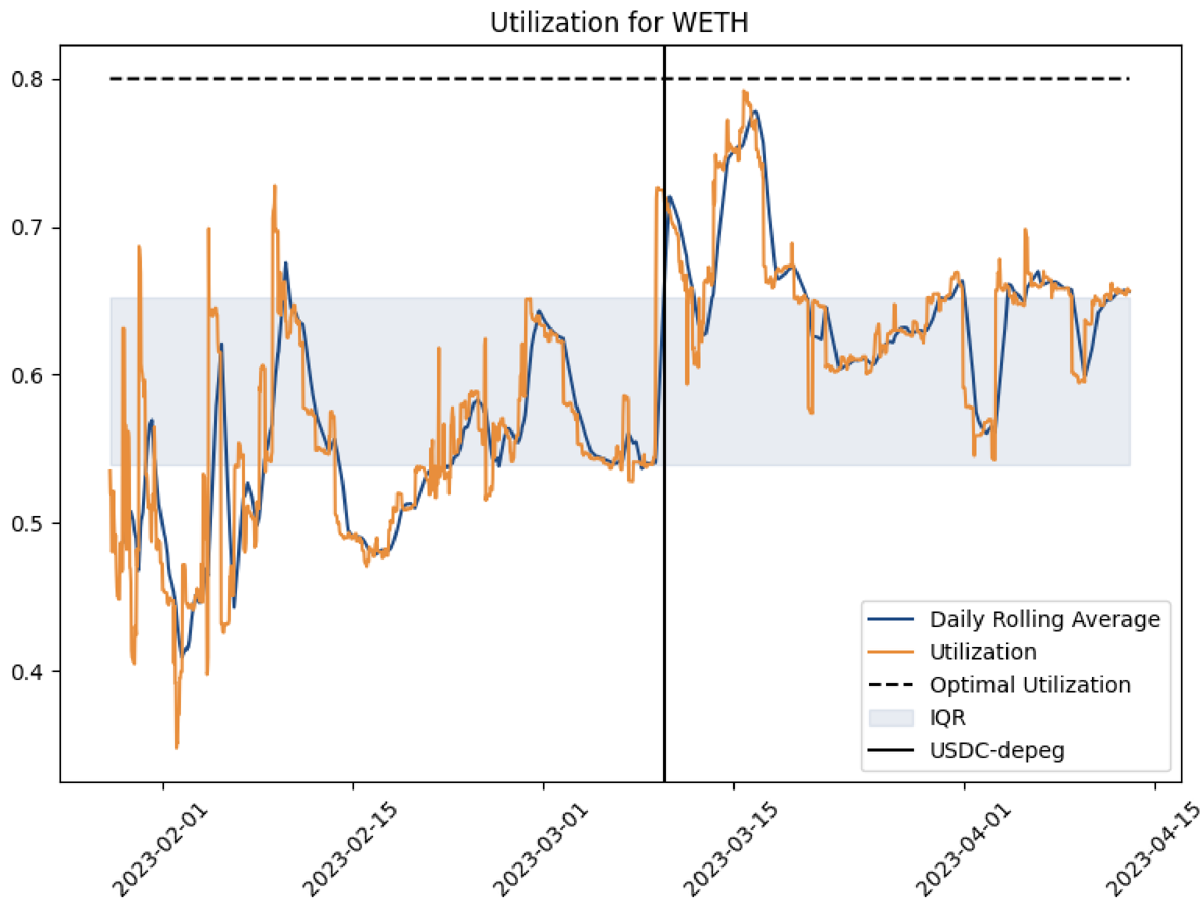

Utilization for WETH
0.8
0.7
0.6
0.5
0.4
Daily Rolling Average
Utilization
Optimal Utilization
IQR
USDC-depeg
2023-02-01
2023-02-15
2023-03-01
2023-03-15
2023-04-01
2023-04-15

Figure 4.11

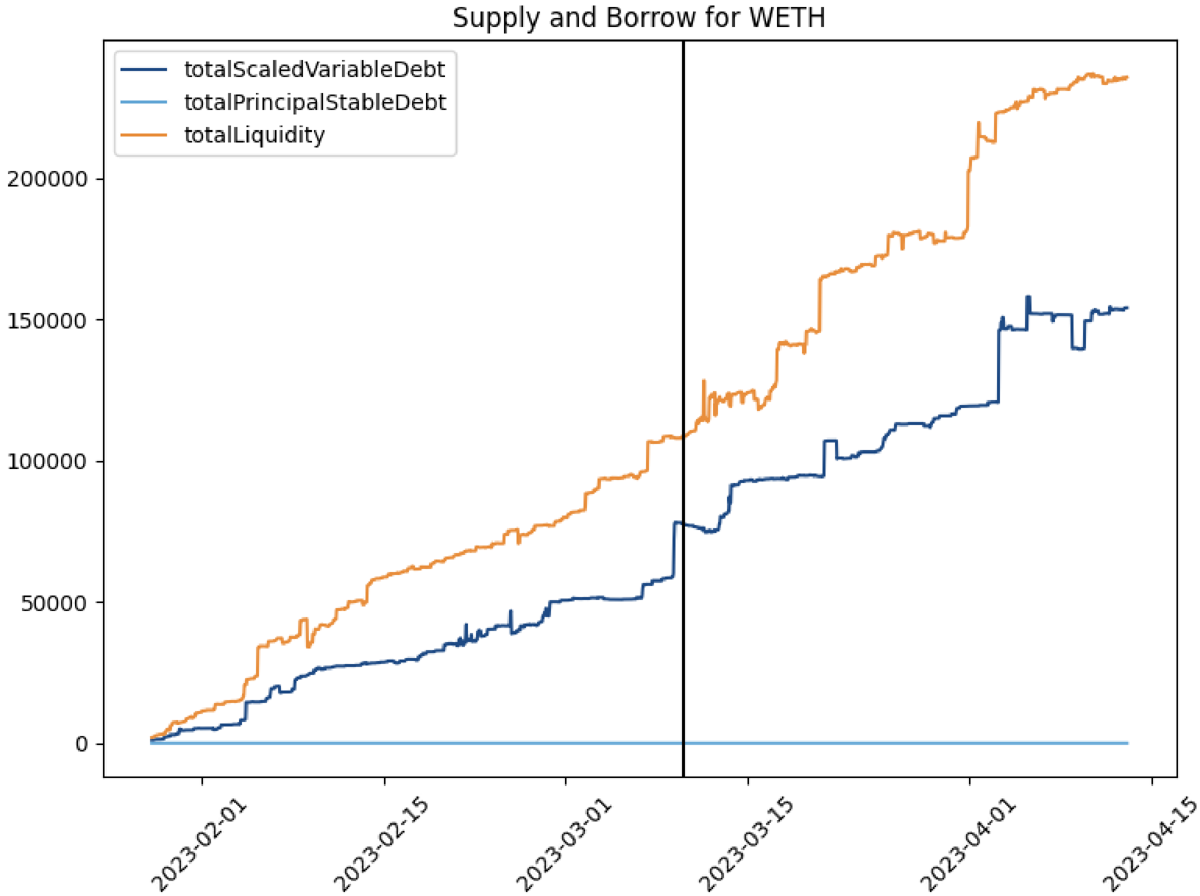
Supply and Borrow for WETH
totalScaledVariableDebt
totalPrincipalStableDebt
totalLiquidity
200000
150000
100000
50000
0
2023-02-01
2023-02-15
2023-03-01
2023-03-15
2023-04-01
2023-04-15

Figure 4.12

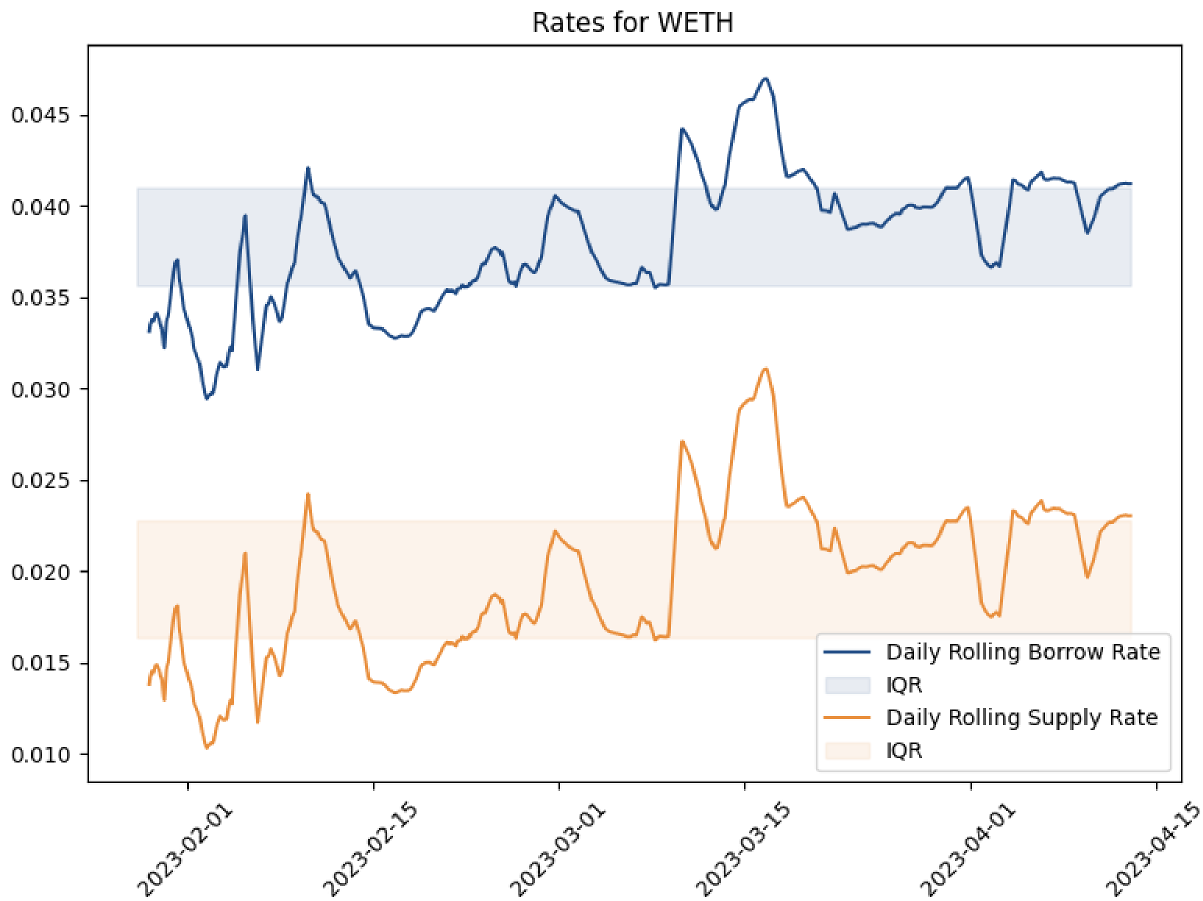

Rates for WETH
0.045
0.040
0.035
0.030
0.025
0.020
0.015
0.010
2023-02-01
2023-02-15
2023-03-01
2023-03-15
2023-04-01
2023-04-15
Daily Rolling Borrow Rate
IQR
Daily Rolling Supply Rate
IQR

Figure 4.13

The utilization rate is significantly higher on v3 compared to v2, with an interquartile range (IRQ) from 54.3% to 65.3% spanning February 1st, 2023, to April 13th, 2023, with a standard deviation $\sigma_u = 8\%$. This yields a target utilization $u_T = 72\%$. Given the pool's proximity to optimal utilization—about one standard deviation off, with a general trend of increasing utilization—we do not propose any modifications to WETH on Aave v3 Ethereum.

3 USDT on Ethereum V2

USDT ranks among the highest-utilization tokens on Aave v2 and has consistently hovered slightly below the optimal rate, showcasing elastic borrowers and suppliers. Consequently, we do not propose any parameter changes to USDT on Aave v2. However, it is worth noting that USDT utilization displays high volatility, with a standard deviation of approximately 10% and a recent prolonged period above optimal utilization following USDC's depeg in March 2023. A potential enhancement for the USDT market could be maintaining the current equilibrium near optimal (70%) while flattening the rate curve to minimize fluctuations.

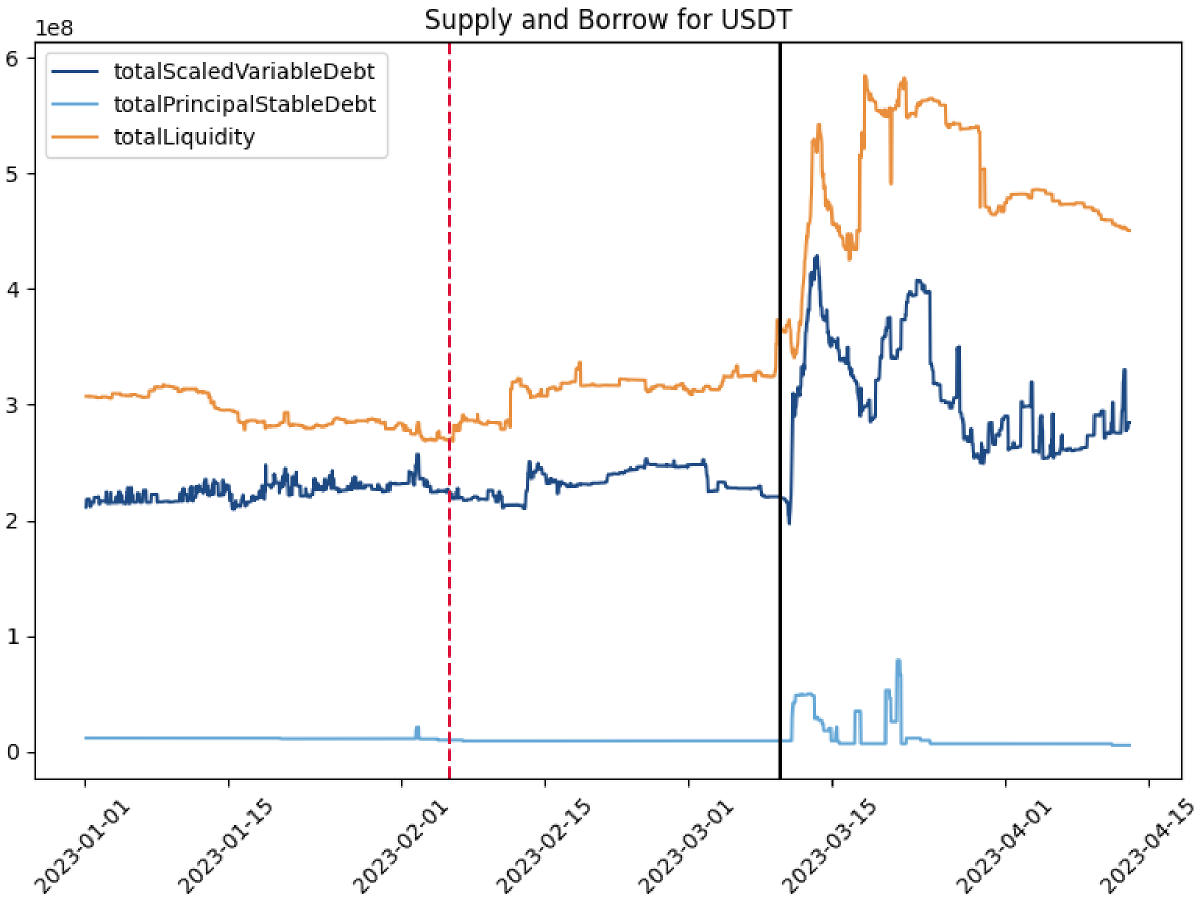
Supply and Borrow for USDT
1e8
6
5
4
3
2
1
0
totalScaledVariableDebt
totalPrincipalStableDebt
totalLiquidity
2023-01-01
2023-01-15
2023-02-01
2023-02-15
2023-03-01
2023-03-15
2023-04-01
2023-04-15

Figure 4.14

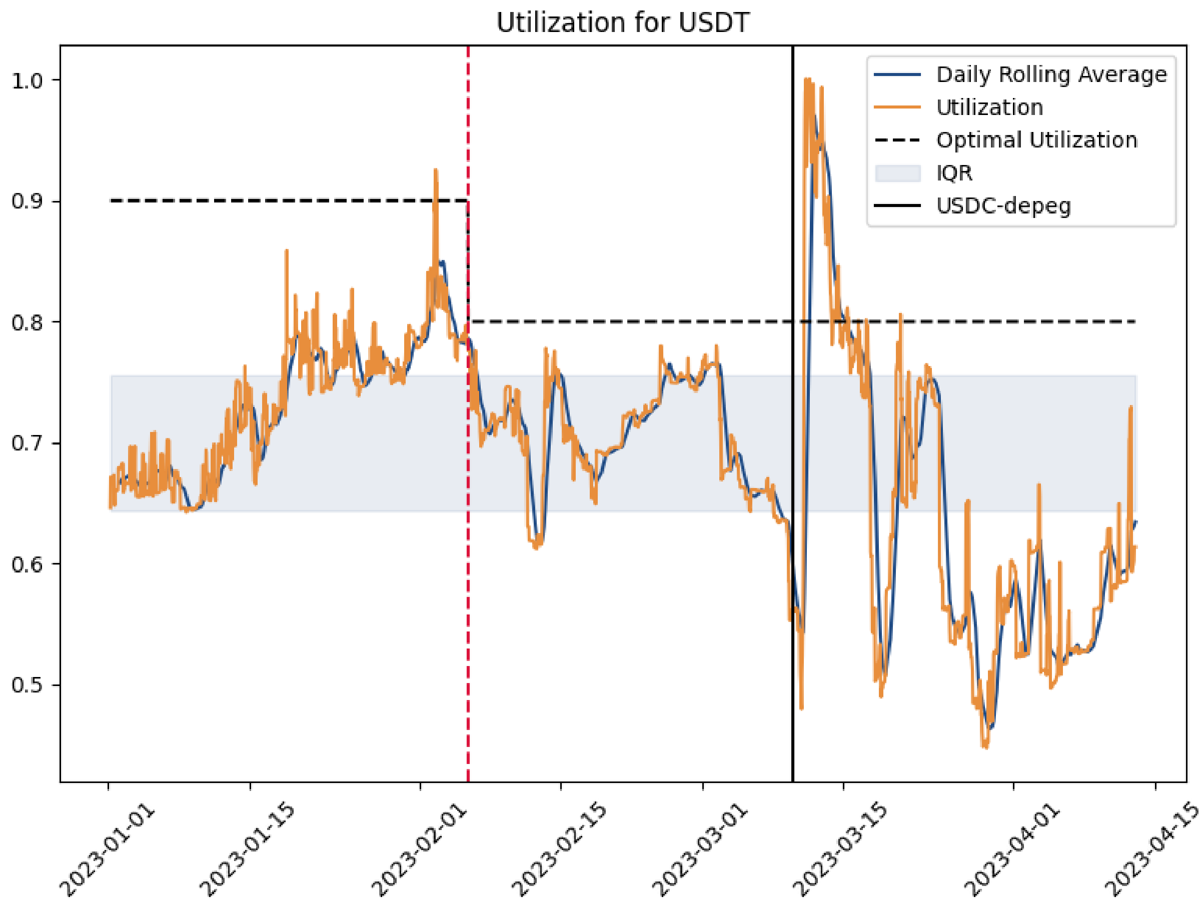
Utilization for USDT
1.0
0.9
0.8
0.7
0.6
0.5
Daily Rolling Average
Utilization
Optimal Utilization
IQR
USDC-depeg
2023-01-01
2023-01-15
2023-02-01
2023-02-15
2023-03-01
2023-03-15
2023-04-01
2023-04-15

Figure 4.15

4 USDT on Ethereum v3

Likewise, USDT's median utilization from February 2023 to April 2023 approximated the target utilization. Consequently, we propose no changes to USDT at this time:

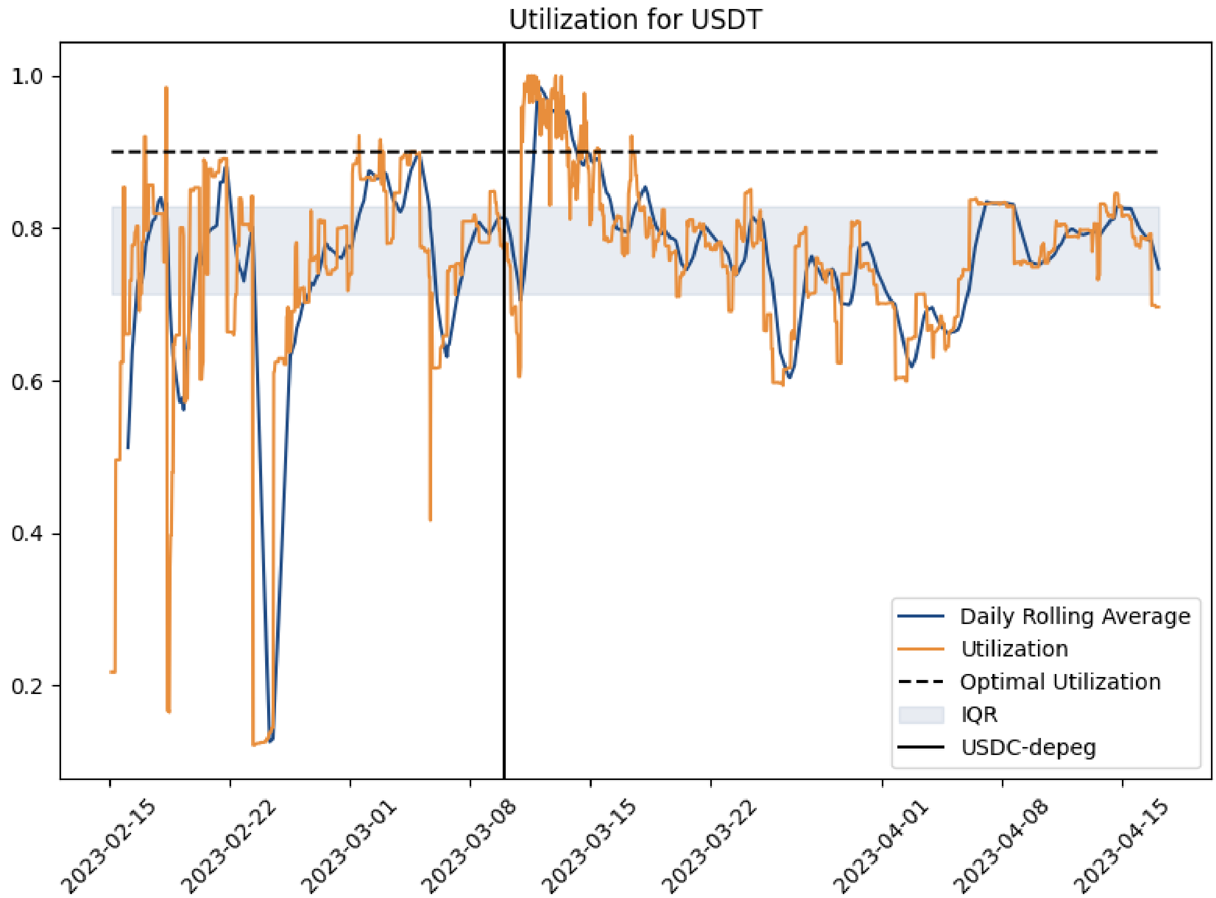
Utilization for USDT
1.0
0.8
0.6
0.4
0.2
Daily Rolling Average
Utilization
Optimal Utilization
IQR
USDC-depeg
2023-02-15
2023-02-22
2023-03-01
2023-03-08
2023-03-15
2023-03-22
2023-04-01
2023-04-08
2023-04-15

Figure 4.16

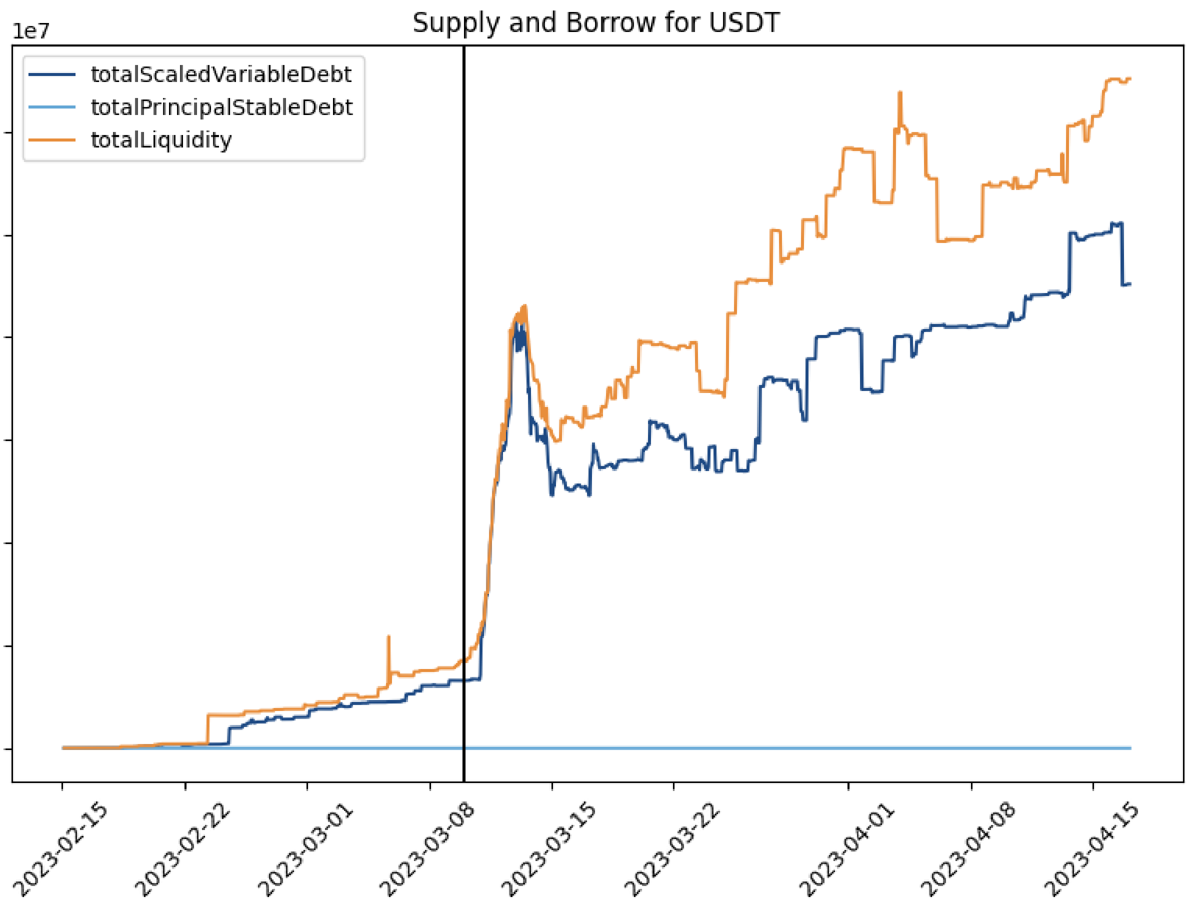
Supply and Borrow for USDT
1e7
totalScaledVariableDebt
totalPrincipalStableDebt
totalLiquidity
0
1
2
3
4
5
6
2023-02-15
2023-02-22
2023-03-01
2023-03-08
2023-03-15
2023-03-22
2023-04-01
2023-04-08
2023-04-15

Figure 4.17

5 Parameter Setting Framework

This section outlines the rhythm at which we anticipate proposing parameter updates and discusses additional aspects relevant to our interest rate-setting strategy.

5.1 New Listings

We will need to frequently propose parameters for tokens that are yet to be listed. In such cases, we suggest a method based on similar assets. For instance, when listing a new low slippage token (LST) on Aave v3's Ethereum platform, we set the optimal utilization according to the lowest existing LST optimal utilization.

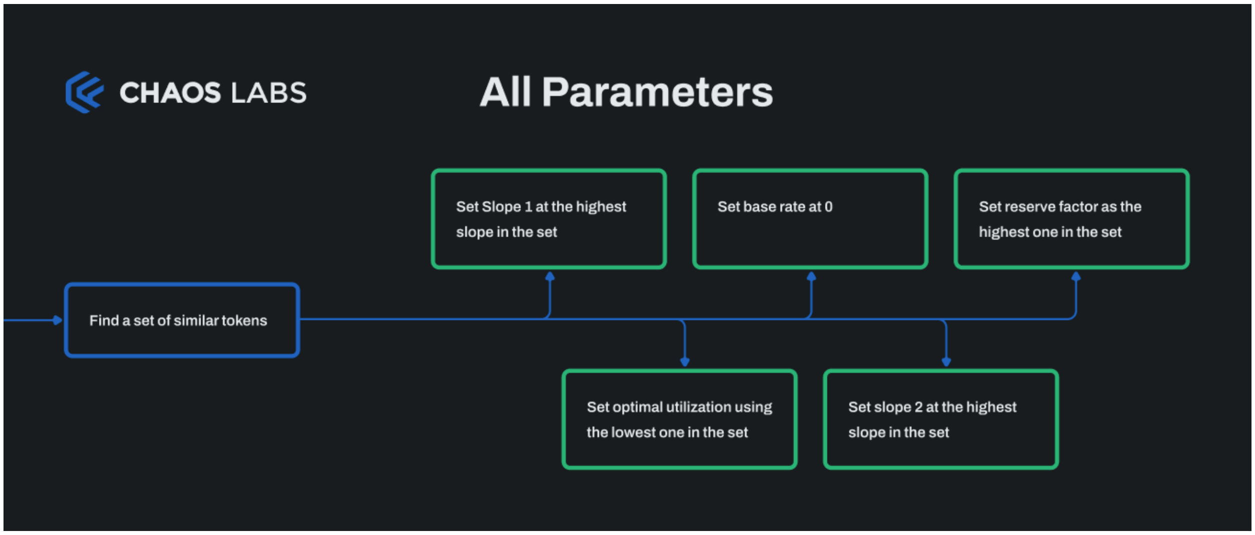

Figure 4.18

Implicitly, we assume one of the following:

1. Tokens in this group share risk profiles, implying similar collateral at risk as a percentage of supply.

2. Tokens in this group experience similar deposit outflows, meaning users withdraw and deposit into the pool at comparable frequencies.

3. Tokens in this group exhibit similar equilibrium interest rates, suggesting suppliers demand analogous APYs.

Following the listing of the asset, we can continuously gather data and perform regular updates if we observe material differences between the set optimal utilization and the newly measured data on collateral at risk or deposit outflows.

5.2 Deployment and Market Migration

Alternatively, if the token is already listed with parameters aligned with our proposed methodology in a separate market or deployment, we might base our initial parameters on these existing parameters. For instance, if we list UNI on a new deployment (e.g., on Binance Smart Chain), we might set UNI's interest rate parameters on BSC according to the average parameters from the existing deployments/markets where UNI is already listed.

5.3 Update Frequency

It is worth noting that risk parameters require frequent recalibration. Hence, we strive to build a sensible framework for setting interest rate curves that we can iterate upon as we accumulate more data and insights into user behavior, and as the market continues to evolve. Importantly, rebalancing interest rate parameters does not immediately trigger liquidations. Thus, establishing a framework for active monitoring allows us to consistently aim for optimal parameters. According to our methodology, we set a lookback period of 3-6 months on our data. Over-utilization is monitored on a monthly basis since such situations warrant immediate attention and entail targeting lower reserve factors and higher borrowing costs. Conversely, to reduce the operational burden on the protocol's governance and service providers, under-utilization is monitored on a quarterly basis.

5.4 Update Condition

Given the operational costs to governance, risk providers, and software/audit providers, we seek to moderate the frequency of parameter updates. We only propose changes to interest rate parameters if the previous quarter's median utilization was one or more standard deviations away from the target.

Example 1: Under-utilization - Optimal utilization $u_{\text{opt}} = 80\%$ - 3-month utilization standard deviation $\sigma_u = 5\%$ - Target utilization $u_T = u_{\text{opt}} - \sigma_u = 75\%$ - 3-month utilization median $u_m = 50\%$ - Conclusion: $u_m < u_T - \sigma_u$. The module is under-utilized and could benefit from a change in parameters.

5.5 V2 to V3 Transition

The transitioning process from v2 to v3 is a crucial aspect we're considering, but our recommendations and methodology do not assume an attempt to deprecate the use of v2. Instead, we strive to target optimal utilization on both v2 and v3. Nevertheless, there is a strong desire from various community members to transition v2 usage to v3, which might lead to targeting a decrease in v2 borrows in favor of v3 borrows. This transition could be achieved by significantly increasing the reserve factor, as often suggested in our methodology, or by increasing the slope to discourage borrowing. A more comprehensive discussion on this transition is outside the scope of this document.

5.6 Stable Borrowing

Stable borrowing is an aspect we have yet to consider in this paper explicitly. Stable borrowing represents a relatively minor portion of all borrowing and can be specifically treated in future studies. The data below illustrates the stable share of all borrowing from January 1st, 2023, to April 1st, 2023, on Aave v2 Ethereum.

Asset	**Borrowing IQR1**	Borrowing IQR2	Borrowing IQR3
1INCH	0.0%	0.0%	0.0%
AAVE	n/a	n/a	n/a
CRV	0.0%	0.0%	0.0%
DAI	1.138%	1.223%	1.269%
ENS	0.0%	0.0%	0.0%
FRAX	0.0%	0.0%	0.0%
GUSD	0.0%	0.0%	0.0%
LINK	1.202%	1.239%	1.239%
LUSD	0.042%	4.195%	4.917%
MKR	0.611%	0.732%	0.844%
SNX	0.0%	0.0%	0.0%
TUSD	14.032%	14.314%	14.882%
UNI	0.001%	0.001%	0.001%
USDC	1.207%	1.261%	1.39%
USDT	3.204%	3.629%	4.289%
WBTC	0.041%	0.254%	0.281%
WETH	0.02%	0.038%	0.054%

5.7 Can we have different parameters for the same token across chains/deployments?

Market rates will likely be different across different chains and deployments. For example, WETH on Avalanche sits at roughly 2.5% borrow rates with 17% in utilization (as of April 17th, 2023), whereas WETH on Ethereum (Aave v3) sits at a 4% borrow rate and 60% utilization. Chains offer different yield and investment opportunities and pose different risks and transaction costs. It follows that interest rate strategies must treat each market and deployment independently.

However, if the borrowing rate at optimal utilization is lower on one chain/deployment than the supply rate at optimal utilization on another chain/deployment, we would expect to observe some arbitrage - although this is not strictly necessary (see WETH borrow and supply rates on Ethereum vs Optimism).

To avoid having arbitrageurs pushing utilization past optimal, we aim to avoid optimal borrow rates being lower on one market/deployment than optimal supply rates on another market/deployment. Effectively:

$$\min(r_{b,m,d}) > \max(r_{s,m,d}), \forall m \in \text{markets}, d \in \text{deployments}$$

5.8 Conclusion

This paper presents a comprehensive framework for refining interest rate strategies in the Aave protocol. The initial phase of the protocol relied on a pre-established strategy due to a lack of data and advanced risk management tools. However, with the availability of empirical data on equilibrium borrow and supply rates, it becomes crucial to consider regular adjustments to interest rate parameters. External factors, such as staking yield, can significantly impact market rates, making adopting a systematic approach for optimizing interest rate curves necessary. This iterative method improves the protocol's reserves and encourages borrowing activity by aiming for optimal utilization and enhancing capital efficiency. As more parameter changes are executed, we will continue to gather data on borrower and supplier responsiveness to changes and formulate expectations on how net borrows and net supplies change when interest rate curves change.

Appendix A

About Chaos Labs

Chaos Labs is a cloud-based platform that develops risk management and economic security tools for decentralized finance (DeFi) protocols. The platform leverages sophisticated and scalable simulations to stress test protocols in adverse and turbulent market conditions. By partnering with DeFi protocols, Chaos Labs aims to create innovative solutions that enhance the efficiency of DeFi marketplaces.

The Chaos Labs team exhibits exceptional talent and represents diverse expertise, encompassing esteemed researchers, engineers, and security professionals. Chaos Labs has garnered its experience and skills from renowned organizations, including Google, Meta, Goldman Sachs, Instagram, Apple, Amazon, and Microsoft. Additionally, the team boasts members who have served in esteemed cyber-intelligence and security military units, further contributing to their unparalleled capabilities.

You can explore our past and ongoing projects for customers like Aave, GMX, Benqi, dYdX, Uniswap, Maker, and more in the Research and Blog sections of our website.

Appendix B

Acknowledgements

We wish to extend our profound appreciation to Professor Dan Galai, whose constructive critiques have been instrumental in shaping this research.