BANK OF CANADA
BANQUE DU CANADA

Staff analytical paper / Document analytique du personnel—2026-13

Last updated: April 2, 2026

DeFi Lending: Returns, Leverage, and Liquidation Risk

Jonathan Chiu
Banking and Payments Department
Bank of Canada
jchiu@bankofcanada.ca

Furkan Danisman
University of Toronto
furkan.danisman@mail.utoronto.ca

DOI: https://doi.org/10.34989/sap-2026-13 | ISSN 3110-9195

DeFi Lending: Returns, Leverage, and Liquidation Risk

Jonathan Chiu[*] Furkan Danisman

Bank of Canada University of Toronto

February 2026

Abstract

We study decentralized lending on Aave V3, the largest DeFi lending protocol by total value locked. Using transaction-level data, we analyze its revenue model, borrower behavior, and liquidation dynamics. We find that protocol earnings are concentrated in a few tokens, that many users engage in recursive leverage despite overcollateralization requirements, and that liquidations occur in concentrated waves but have limited impacts on broader markets. Overall, DeFi lending with proper governance is operationally viable, but it also faces constraints related to capital efficiency, liquidation risk, and systemic fragility within the crypto ecosystem.

JEL classification: E50, E58

Keywords: Blockchains, Decentralized Finance, Liquidation, Smart Contracts

[*]We thank many friends and colleagues for many useful conversations. The views expressed here are those of the authors and do not necessarily reflect the position of the Bank of Canada.

1 Introduction

Decentralized finance (DeFi) refers to blockchain-based protocols that facilitate financial transactions—such as lending, borrowing, and asset exchange—without centralized intermediation or oversight. Among the various DeFi applications, decentralized lending protocols have attracted particular attention due to their growing scale and functional resemblance to traditional financial intermediation.[1] As of 2025, lending protocols such as Aave account for a substantial share of total value locked (TVL) in DeFi, facilitating billions of dollars in DeFi lending activity.[2] Although DeFi lending protocols depart from traditional banks in terms of regulation, identity verification of customers, and risk management, they perform similar credit functions by using smart contracts, the distinctive technology of DeFi, alongside collateralized borrowing. Understanding how DeFi lending works—how it allocates capital, manages risk, and affects user behavior—is important for assessing its potential benefits and risks, both for researchers and for regulators.

This paper provides a detailed empirical examination of decentralized lending using transaction-level data from Aave Version 3 (V3), the largest DeFi lending protocol by TVL. We begin by contrasting the institutional mechanics of DeFi lending with those of traditional banking, highlighting key differences in regulation, identity verification, intermediation, and risk management. We then focus on three aspects of the DeFi lending ecosystem: (i) the protocol's revenue model, (ii) the prevalence of margin trading activity despite strict overcollateralization requirements, and (iii) the triggers, dynamics and consequences of liquidations. These are understudied in the existing literature, yet they are central because they highlight the key channels through which DeFi lending affects efficiency (capital allocation, pricing, and participation) and fragility (liquidation cascades, liquidity shortfalls, and shock propagation), and therefore the systemic implications of decentralized credit markets.

First, we examine how the platform generates revenue and profits. Unlike traditional banks, which operate under high fixed costs and rely on interest margins to fund operations and manage credit risk, DeFi lending protocols such as Aave can be built with relatively low technological and oper-

[1]Below, we use the terms "protocol" and "platform" interchangeably, since in DeFi the user-facing platform is typically governed by the underlying smart-contract protocol.

[2]TVL is the current market value of all crypto assets that users have deposited into a DeFi protocol's smart contracts so that those assets are locked in the protocol at a point in time. It is often used as a snapshot measure of scale.

ational costs through smart contracts (though they may incur substantial initial costs to attract users and liquidity). To assess how these features translate into actual performance, we quantify the protocol's expected and realized earnings and disaggregate profitability by token. While a lending platform typically facilitates the lending of many different tokens, we show that a small subset of tokens accounts for most of the platform's lending revenue.

Second, we investigate the practice of margin trading via *recursive leverage*, whereby a user repeatedly borrows against collateral, swaps the borrowed funds for additional collateral, and re-deposits it into the protocol to borrow again. Despite the requirement for overcollateralization, many users engage in borrowing not to obtain liquidity for external spending, but to amplify exposure to certain assets. We develop a transaction-based methodology to identify such behavior and find that margin trading accounts for a significant share of borrowing activity, particularly among large, sophisticated users. These traders operate with greater frequency and size, and are significantly more exposed to liquidation risk.

Third, we analyze liquidation patterns as both a risk control mechanism and a source of stress. Liquidations are essential to maintaining solvency in the absence of identity-based enforcement. We document that liquidations on Aave V3 occur in clustered waves, driven primarily by sharp declines in collateral prices. We show that large borrowers are disproportionately affected during peak events and estimate that realized losses—including liquidation penalties and missed price recoveries—can amount to 10–30% of liquidated value.

Together, these findings suggest that, while DeFi lending provides transparency, automation, and low-cost execution, it relies on strict collateral requirements and exposes users to volatile and sometimes abrupt adjustments. DeFi lending mitigates some intermediation frictions—such as costly monitoring and slow execution—but does so by limiting credit flexibility through overcollateralization and rule-based liquidation.

The growing literature on decentralized lending includes both empirical and theoretical contributions. Empirical studies such as [Qin et al., 2021], [Perez et al., 2020], [Heimbach et al., 2023b], [Heimbach and Huang, 2024], and [Lehar and Parlour, 2022] examine user behavior, liquidation dynamics, leverage patterns, and potential spillovers in DeFi lending protocols. These papers provide evidence on how platform design and user incentives interact to shape systemic outcomes.

On the theoretical side, [Kozhan et al., 2024], [Rivera et al., 2023], and [Chiu et al., 2023] offer models that analyze recursive leverage, interest rate design, and fragility in DeFi credit markets. Together, this body of work highlights both the operational viability and the inherent constraints of lending without intermediation. Our paper contributes to this literature by combining transaction-level analysis with protocol-level measures of profitability, leverage, and liquidation outcomes.

The remainder of the paper is organized as follows. Section 2 provides institutional background on Aave V3 and compares it with traditional bank lending. Section 3 analyzes the protocol's revenue model and token-level performance. Section 4 investigates the prevalence and characteristics of margin trading among borrowers. Section 5 studies liquidation patterns and their implications for user outcomes and market stability. Section 6 concludes.

2 DeFi Lending: Institutional Information and Data Sources

2.1 Aave Lending Protocol

We focus on Aave V3 because it is the largest decentralized lending protocol by total value locked (TVL). According to DefiLlama,[3] Aave is also the largest decentralized finance protocol overall, with approximately 34 billion USD secured in smart contracts. This accounts for nearly 25% of the TVL across all decentralized protocols and 50% of the TVL within the lending sector.

Aave V3 is a decentralized, non-custodial liquidity protocol that enables users to lend and borrow crypto assets through smart contracts on the blockchain. Aave is deployed across multiple blockchain networks, including Ethereum, Polygon, Avalanche, Arbitrum, Optimism, Base and BNB Chain. Examples of assets accepted by Aave include stablecoins such as USDC, USDT, DAI, and GHO, ETH-linked assets such as ETH, WETH, and wstETH, and BTC-linked assets such as WBTC and cbBTC. Users supply assets into liquidity pools and receive aTokens in return. For example, when a user supplies an asset WETH, Aave mints an aToken aWETH; when the user withdraws, the user redeems that aToken for the underlying asset, including accrued interest. These aTokens are interest-bearing and pegged 1:1 to the underlying asset. Interest accrued is automatically reflected in the user's aToken balance, which, normally, can be transferred or redeemed at any time. Note that when users interact directly with Aave, they must pay blockchain gas fees for on-chain transactions such as supplying, borrowing, repaying, or withdrawing

[3] Accessed on July 29, 2025. Available at: `https://defillama.com/`

Borrowers can draw from these pooled funds by locking up collateral. All loans are overcollateralized, meaning the value of the collateral must exceed the borrowed amount. Aave V3 supports both variable and stable interest rate options. Interest rates are determined algorithmically based on the utilization rate of each asset pool. As utilization increases—indicating higher borrowing demand—interest rates adjust upward to incentivize additional supply and moderate borrowing. Aave V3 introduces several improvements over previous versions, including isolation mode, efficiency mode, and siloed borrowing.

Each collateral asset has associated risk parameters, including a loan-to-value (LTV) ratio and a liquidation threshold. The LTV determines how much can be borrowed against a given collateral, while the liquidation threshold defines the point at which the position becomes undercollateralized. A measure called "health factor" is used to determine if a position is undercollateralized. If the health factor of a user's account drops below one (typically due to falling collateral prices or accrued interest), the position becomes eligible for liquidation. Third-party liquidators can repay a portion of the debt and receive the corresponding collateral, plus a liquidation bonus.

To protect against protocol-level losses, Aave employs a decentralized insurance mechanism called the Safety Module. Aave token holders can stake their tokens to provide a capital buffer. In the event of a shortfall event—such as widespread liquidation failures or oracle manipulation-up to 30% of the staked Aave can be slashed to cover bad debt. Stakers are compensated with protocol fees and safety incentives. If the Safety Module cannot fully cover the losses, an emergency recovery issuance of new Aave tokens can be triggered to recapitalize the protocol.

Aave V3 also supports cross-chain functionality via the Portal feature, which allows assets to be moved across supported blockchains under the same protocol, enabling unified liquidity across chains. Governance remains decentralized and community-driven, with Aave token holders voting on key protocol parameters, asset listings, and upgrades.

2.2 Comparison with Traditional Bank Lending

The institutional design of decentralized finance lending protocols like Aave differs fundamentally from that of traditional financial institutions. As shown in Table 1, these differences shape how credit is extended, risks are managed, and trust is established.

First, **regulation** forms a core pillar of traditional banking. Banks operate under extensive pru-

Feature	**Bank Lending**	**DeFi Lending**
Regulations	Reserve, liquidity, capital requirements, etc.	Only rules defined by smart contracts
Borrower	Identified	Anonymous
Collateral	Standard assets	Crypto assets
Intermediary	Trusted human actors	Pre-programmed smart contracts

Table 1: Institutional differences between Bank and DeFi lending

dential oversight, including capital adequacy, liquidity coverage, deposit insurance, and reserve requirements. In contrast, DeFi protocols are permissionless and operate in a largely unregulated space, with no formal supervisory oversight or mandated risk buffers. Risk management is instead embedded in protocol design-primarily through overcollateralization, automated liquidations, and open-source smart contracts.

Second, **borrower identity** plays a critical role in traditional lending. Banks verify the identity, creditworthiness, and financial history of borrowers before extending loans. Because borrowers are known and subject to legal recourse, repayment can be enforced through reputation, relationship banking, and the threat of litigation. By contrast, DeFi lending is *pseudonymous*: users interact with the protocol through blockchain addresses without undergoing identity verification. As a result, DeFi protocols cannot rely on trust or reputation; enforcement is purely *collateral-based*, and all loans must be overcollateralized to mitigate default risk.

Third, the **collateral base** in DeFi is fundamentally different. Traditional banks typically accept standard financial or physical assets—such as real estate, securities, or guaranteed income streams—as collateral. In DeFi, typically only *on-chain crypto assets* are accepted. These assets tend to be more volatile and less correlated with traditional macroeconomic factors, requiring higher haircuts and more frequent revaluation.

Finally, the nature of **intermediation** diverges sharply. Traditional banks rely on *trusted human actors*—bank officers, analysts, and risk managers—who exercise judgment in underwriting, monitoring, and restructuring loans. DeFi protocols replace human discretion with *pre-programmed smart contracts*, which automatically enforce lending terms, interest accrual, and liquidation rules without subjective input. While this automation reduces costs and eliminates counterparty discretion, it also introduces rigidity and new forms of technical risk.

In sum, while both systems aim to facilitate credit intermediation, they rely on fundamentally different mechanisms to establish trust, enforce contracts, and manage risk.

2.3 Comparison with Traditional Banking

To better understand the distinctiveness of decentralized lending, it is instructive to compare Aave V3 to traditional banks using key balance sheet and performance metrics. Table 2 summarizes total loans, the loan-to-deposit ratio, net interest margin (NIM), and non-performing loans (NPL) for the five major U.S. banks, five major Canadian banks, and Aave V3 on Ethereum in 2024. For the U.S. banks, we consider JPMorgan Chase, Bank of America, Citigroup, Wells Fargo, and U.S. Bancorp. For the Canadian banks, we consider RBC, TD Bank, Scotiabank, BMO, and CIBC.

The data for the U.S. and Canadian banks are drawn from their publicly available 2024 annual reports. To clarify terminology differences between traditional and decentralized finance, in the DeFi context, the term "supply" functions as an analog to "deposits". Throughout this paper, we use supply and deposit interchangeably. Likewise, the traditional loan-to-deposit ratio maps directly to the utilization rate in Aave.

In conventional banking, the net interest margin (NIM) is calculated as net interest income divided by average earning assets. For Aave V3, we compute an analogous NIM as the total net income from lending operations divided by the aggregate borrowed amount in USD. Finally, the DeFi equivalent to non-performing loans (NPL) is the bad debt, defined as unrecovered borrowing that remains after the borrower's collateral has been fully liquidated.

Having established comparable financial metrics, we now contrast the scale and structure of Aave V3 with those of traditional U.S. and Canadian banks. Aave V3 operates at a much smaller scale than its traditional counterparts. While the major U.S. and Canadian banks in Table 2 reported outstanding loans of over 800 billion USD and 500 billion USD, respectively, Aave V3's total lending volume in 2024 amounted to roughly 6.0 billion USD. This reflects not only differences in institutional maturity and market scope but also Aave's design constraint that all loans must be overcollateralized and backed by on-chain crypto assets.

The loan-to-deposit ratio for Aave V3 was 40%, significantly lower than the average ratios of 61.2% and 74.2% for the major U.S. and Canadian banks, respectively. This conservative ratio arises from Aave's structural separation of supplied and borrowable assets across risk pools and the protocol's overcollateralization requirement. In contrast, banks are able to lend a larger fraction of their deposits due to access to credit histories, regulatory capital buffers, and deposit insurance.

Another striking difference lies in the net interest margin (NIM). The average NIM was 2.48% for the major U.S. banks and 1.69% for the major Canadian banks, compared with only 0.64% for Aave V3. Part of this difference reflects the inherently lower risk in Aave's model, which avoids unsecured lending. However, a key reason for banks' higher spread is their greater intermediation cost: salaries, branch infrastructure, compliance, and customer support services must all be funded from the interest margin. In contrast, Aave operates on minimal overhead via smart contracts, enabling lower spreads—but also providing fewer services and less flexibility to borrowers.

In terms of credit performance, Aave V3 reported zero non-performing loans (NPL), while the average NPL ratio was 0.59% for the major U.S. banks and 0.65% for the major Canadian banks. This difference is due to Aave's automated and real-time liquidation system, which preempts defaults by enforcing collateral adequacy through smart contracts. Deposits are never exposed to undercollateralized risk, but borrowers bear the downside of sharp price movements in volatile collateral.

Taken together, these comparisons highlight the contrasting design philosophies of the two systems. Banks actively manage credit risk and customer relationships under regulation, whereas Aave provides an automated, transparent platform with minimal discretionary control. While DeFi's architecture limits intermediation capacity and credit flexibility, it enables continuous operation with lower cost, no counterparty discretion, and no credit defaults—at the price of strict collateral requirements and price-based enforcement.

Platform	Loan	$\frac{Loan}{Deposit}$	NIM	NPL
Major U.S. Banks	886,178	61.2%	2.48%	0.59%
Major Canadian Banks	569,247	74.2%	1.69%	0.65%
Aave V3	6,024	40.0%	0.64%	0.00%

Table 2: Summary of average total loans, loan-to-deposit ratio, net interest margin (NIM), and non-performing loans (NPL) for the five major U.S. and five major Canadian banks, and Aave V3 on Ethereum in 2024. All monetary values are in millions of USD.

2.4 Data Description

This study focuses exclusively on the Ethereum blockchain to examine platform and user level activities on Aave V3. The sample spans from January 27, 2023, the launch date of Version 3, through May 6, 2025. While we acknowledge that this scope does not include other versions or blockchains supported by Aave, Ethereum accounts for approximately 85% of Aave's total value locked, and Version 3 on average represents 72% of total deposits across all protocol versions during the sample period. We therefore consider this focus both representative and analytically tractable. Furthermore, we note that activity on Version 2 declines to 1% of total deposits during the sample period, while Version 1 was decommissioned prior to our observation window.

The primary dataset is constructed using publicly available data from Dune Analytics [4], which provides structured, curated blockchain snapshots and historical event logs. To supplement deeper user-level activities, we extract additional information from Etherscan [5] for selected addresses of interest. Token market capitalizations, used in our analysis of liquidation impacts, are retrieved from Coingecko [6], while the CCI30 index [7] is employed as a proxy for overall crypto market performance.

Missing values were assumed to be missing at random (MAR) and were excluded from the analysis accordingly.

3 Platform's Revenue model

In DeFi lending platforms such as Aave, the overcollateralized nature of the protocol ensures that, for each asset, total supplied liquidity consistently exceeds total borrowings. Importantly, the collateral and borrowed tokens can differ across positions. For example, the protocol may have a supply of 100 US dollars worth of both WETH and USDC. Users may use WETH as collateral to borrow USDC, while no borrowing occurs in the opposite direction. Although this structure enhances efficiency across liquidity pools, it also implies that certain tokens generate positive net revenue for the protocol, while others may not. In this section, we evaluate the financial contribution of individual tokens, as well as their aggregated impact, to the platform's performance on Aave V3.

[4] Accessed on July 29, 2025. Available at: `https://dune.com`.
[5] Accessed on July 29, 2025. Available at: `https://etherscan.io`
[6] Accessed on July 29, 2025. Available at: `https://www.coingecko.com`
[7] Accessed on July 29, 2025. Available at: `https://cci30.com`

Each token on the platform is associated with a supply annual percentage yield (APY) and a borrow annual percentage yield (APY). As part of the protocol's revenue generating structure, the borrow APY is systematically set higher than the supply APY. To approximate the protocol's expected return from lending activity, we define a token specific rate of return (RoR) as follows:

$$\mathrm{RoR}^i = \frac{\mathrm{Borrow}^i(\mathrm{APY}) \times \mathrm{UR} - \mathrm{Supply}^i(\mathrm{APY})}{365},$$

where i denotes the token and UR represents the utilization rate of the asset.

Earning from Lending

In addition to the spread between borrowing and lending rates, the protocol may offer incentives to encourage users to supply or borrow specific tokens. These incentives are distributed in the form of additional token rewards. We refer to this component as the market incentive (MI). To obtain a more accurate estimate of the net earning, we subtract the value of these incentives from the return. Accordingly, we define the earning from lending (EFL) for token i as:

$$\mathrm{EFL}^i = \left(\frac{\mathrm{RoR}^i \times \mathrm{Supply}^i}{100} - \mathrm{MI}^i\right).$$

This formulation estimates the platform's earning per token, excluding operational costs such as infrastructure or staffing. It focuses exclusively on the revenue and expenses intrinsic to the lending market. However, it is important to note that these earnings are not realized until the underlying loans are either settled or partially repaid by the borrower.

Figure 1 presents the aggregated daily earning from lending derived from lending activities, alongside the total circulating supply and borrowing volume on V3. Circulating supply and borrow are computed as the cumulative net difference between tokens minted and tokens burned over time. During the sample period, the total supplied volume peaked at approximately 30 billion US dollars, while borrowings reached around 12 billion US dollars. The maximum daily earning from lending approached 200 thousand US dollars.

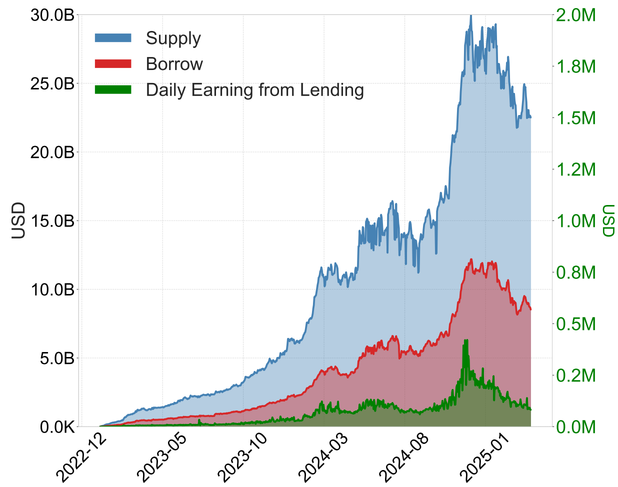

Figure 1: Daily lending earnings, circulating supply, and borrowing volumes (USD) on Aave V3.

Realized Earnings

It is important to note that the protocol's revenue model is not limited to secured lending. Additional income sources include liquidation fees, flash loan fees, and other service fees. To assess realized earnings, we track tokens minted to the protocol's treasury address, which also includes the revenue from lending and subtract any market incentive. Figure 2 presents the seven-day moving average of realized earnings over time for V3. The series exhibits multiple peaks, with daily values occasionally surpassing 600,000 to 800,000 USD.

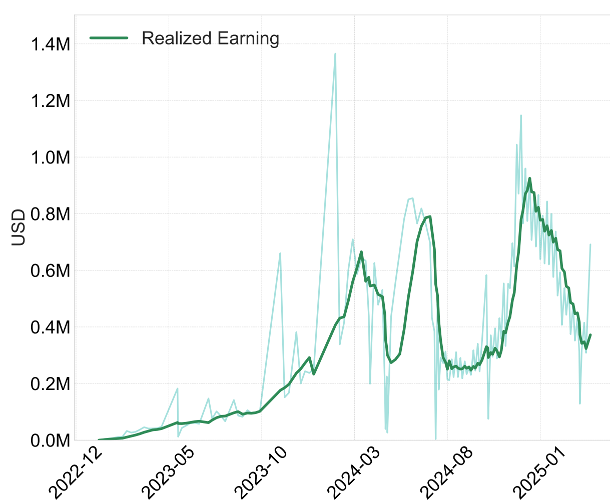

Figure 2: Seven-day moving average of the daily realized earnings (USD) on Aave V3.

Average Rate of Return

The platform's interest rate model is designed to maintain market stability by dynamically adjusting the spread between borrow and supply rates as utilization increases. Figure 3 illustrates the weighted average spread between borrow and supply rates and the rate of return based on supply share. The spread fluctuates around 1 percent over time. In contrast, the protocol's rate of return remains more stable, at approximately 0.25 percent, suggesting that on average the platform earns 0.25 cents per dollar of supplied liquidity.

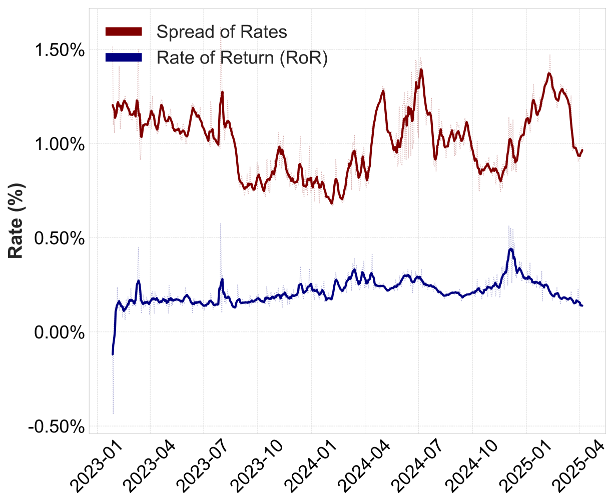

Figure 3: Weighted average spread between borrow and supply rates, and average rate of return over time based on supply share.

Returns and Earnings by Tokens

This aggregated information is further disaggregated at the token level in Table 3. The token with the highest supply share is weETH, accounting for approximately 20 percent of total supply. However, its contribution to the platform's overall earning is relatively limited compared to WETH, USDT, and USDC, which together account for nearly 83 percent of total earning.[8]

From the Table 3, we observe that while all three of these tokens (WETH, USDT, and USDC) exhibit relatively high utilization rates and make comparable contributions to overall earning,

[8]Note that ETH denotes ordinary Ether. WETH is the same asset in a token format used more easily in DeFi, while wstETH, weETH, and rsETH are tokenized claims on staked or restaked ETH that are designed to earn yield and therefore are more complex than plain ETH.

Symbol	Supply(%)	Borrow(%)	UR	Supply APY(%)	Rate of Return(%)	Earning(%)
weETH	**19.70**	2.12	0.05	0.08	0.06	4.03
wstETH	16.71	4.21	0.05	0.03	0.00	0.21
WETH	15.34	43.07	0.76	1.89	0.33	**27.49**
WBTC	10.32	3.47	0.11	0.10	0.02	1.13
USDT	8.03	21.22	0.79	5.40	0.53	**26.40**
USDC	7.33	20.40	**0.83**	5.46	0.63	**28.80**
rsETH	6.85	0.00	0.00	0.00	-0.00	-0.00
cbBTC	4.27	0.53	0.04	0.01	0.00	0.07
sUSDe	2.59	0.00	0.00	0.00	0.00	0.00
USDS	1.61	1.73	0.36	3.94	0.43	3.51

Table 3: Top 10 cryptocurrencies with the highest market share on Aave V3, including their deposit and borrowing shares, utilization rate, supply APY, rate of return, and earning share.

WETH has a lower rate of return than USDT and USDC. This gap is offset by WETH's higher supply share, which is nearly double that of each of the other two assets.

These findings suggest that only a small subset of tokens are responsible for the majority of the platform's earning. Furthermore, a high supply share alone, such as that observed for weETH and wstETH, does not necessarily translate into high earning for the protocol if utilization remains low.

Default Risk and Backstop Mechanism

The lending platform and token suppliers are exposed to default risk if borrowers fail to repay their loans, and the collateral posted is insufficient to cover the resulting shortfall. To mitigate this risk, Aave introduces a backstop mechanism known as the Safety Module, which serves as a decentralized form of insurance. The Safety Module consists of a pool of assets—primarily Aave tokens—that can be used to cover protocol shortfalls arising from insolvency events. This reserve is funded both by the protocol and by users who voluntarily stake tokens into the Safety Module. In return, these users are incentivized through reward yields and governance rights.

Effective stake on the insurance per Aave V3 Supply unit is plotted in Figure 4, measured as the amount of Safety Module backing per dollar supplied. At the beginning of the V3, the supply value is low, hence there were about 10 cents invested into the insurance scheme for each supplied dollar. This amount gradually declined to around 2 cents in 2025, suggesting that backstop protection became thinner relative to the scale of activity as V3 grew, although no backstop event occurred during the sample. Nonetheless, the protocol's insurance framework was previously tested in a high-profile incident on Version 2, where it successfully absorbed losses without impacting user deposits, illustrating the resilience of Aave's backstop mechanisms under stress.[9] In response to such vulnerabilities, Aave introduced additional safeguards in V3 to limit risks and reinforce protocol resilience in the event of market stress.

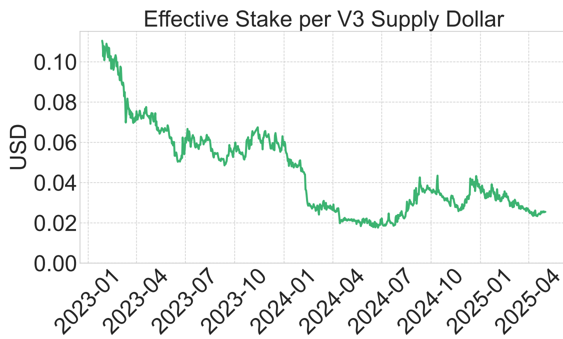

Figure 4: Effective insurance amount per US dollar supplied.

4 Margin Trading Activities

Among the various reasons for borrowers to use DeFi lending protocols like Aave, margin trading through recursive leverage is a particularly interesting one.[10] In traditional finance, margin trading refers to the practice of borrowing funds to trade financial assets, typically using the assets themselves as collateral. A margin account is often required, where the trader must maintain a minimum margin (collateral relative to the borrowed funds).

While borrowing is subject to over-collateralization, recursive leverage on a lending platform could be a way to replicate margin trading behavior. This involves repeated cycles of borrowing, swapping, and collateralizing. For instance, a user may supply token A, borrow token B, swap the borrowed

[9] On November 14, 2022, a trader publicly disclosed a $63 million short position on the Curve (CRV) token using Aave V2. Following a price rebound in CRV, the position was liquidated, resulting in $1.3 million in bad debt. Aave covered the loss using funds from the Safety Module and treasury, as approved by the community through governance. See [Warmuz et al., 2023, Heimbach et al., 2023a]

[10] There are other reasons. For instance, a borrower may want to unlock liquidity from long-term holdings of a crypto asset without triggering a taxable event or relinquishing ownership; A borrower may want to borrow to exploit rate differentials or price inefficiencies across platforms; Borrowers may be motivated by token rewards or fee rebates offered by protocols to subsidize borrowing.

token B for additional token A, and then use the newly acquired token A as collateral to borrow more of token B. This strategy is typically motivated by the expectation that token A will appreciate in value. While this mechanism amplifies potential gains, it also magnifies potential losses. The availability of such recursive behaviour is a unique feature of decentralized finance lending and is not feasible under traditional lending frameworks.

Identification of Margin Trading

To empirically identify this behavior, we define margin trading activity as follows: Suppose a user holds token A and uses it as collateral to borrow token B. If, within one day of the borrowing event, the user swaps token B for additional token A, and subsequently, within the next day, uses the newly acquired token A to borrow more of token B, we classify this sequence of transactions as margin trading borrowing. This identification strategy provides an approximation of margin trading behavior. We also tested alternate definitions based on longer time windows and found no significant difference in the resulting estimates.

Using this definition, we find that margin trading activity accounts for approximately 20 percent of the total borrowed volume and 8 percent of the total number of borrowing transactions, as shown in Table 4. This indicates that a significant share of borrowing activity is driven by users engaged in margin trading strategies and these loans tend to be larger than the average size.

	Borrow Amount ($)	Borrow Number
Margin Trading	22B	25K
Total	103B	300K
Margin Trading (%)	20.46%	8.20%

Table 4: Summary of recursive leverage activity within a 1-day window following initial borrow events. The leverage column reflects the volume and number of borrow events where users swapped borrowed assets and re-borrowed within one day. Total borrow metrics correspond to all borrow activity observed between January 1, 2023, and May 1, 2025.

Time Series Pattern

Figure 5 presents the time series of margin trading activity over the sample period. Following the launch of V3, this form of borrowing behavior rapidly gained popularity, reaching a peak of approximately 40 percent of total borrowing volume within the first six months. Interestingly, while the volume of margin trading increased, the number of such transactions remained relatively stable. This implies that users engaged in margin trading began borrowing larger amounts per transaction, suggesting a rise in the average loan size within this category.

A notable decline in margin trading volume is observed after a major liquidation event in August 2024, during which 258 million US dollars get liquidated, as documented in Table 6. This sharp contraction is not necessarily a voluntary reduction in margin trading by users. Instead, it is possible that the liquidation shock resulted in substantial realized losses for leveraged borrowers, thereby reducing their capacity or willingness to reenter such positions. Following the dissipation of this shock, margin trading activity gradually recovered, both in terms of volume and number of transactions.

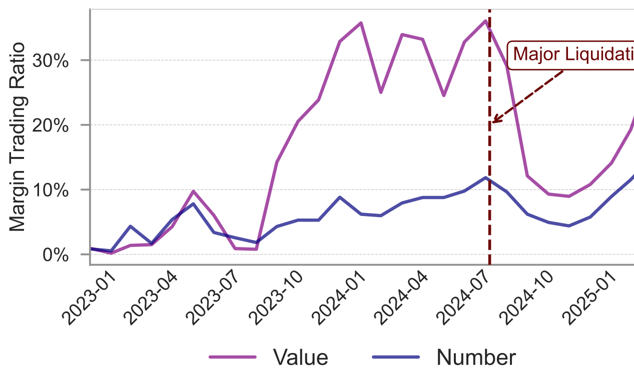

Figure 5: Proportional margin trading borrowing volume and frequency relative to total borrowing between January 1, 2023, and May 1, 2025.

User Level Pattern

We observe that this borrowing behavior accounts for a substantial portion of total borrowing volume and plays an important role in the decentralized finance lending landscape. Users who engage in margin trading tend to exhibit distinct behavioral and financial characteristics. Some descriptive statistics are summarized in Table 5.

Although margin traders account for a large share of total borrowing volume, they represent only a small fraction, approximately 2 percent, of all active users on the platform. Using the classification framework for large investors and retail users from [Park and Stinner, 2023], we find that margin traders are disproportionately composed of large investors (58%). Relative to their counterparts, margin traders are also considerably more "active", which we define as having an outstanding debt position.

Metric	**Non-Margin Traders**	**Margin Traders**
Active user	29,778 (98%)	579 (2%)
Large investor	9782 (33%)	337 (58%)
Retail user	10033 (34%)	86 (15%)
Debt ($)	0.01	0.07
Active weeks	12.0	23.0
Borrow frequency (day)	0.17	0.92
Borrow amount ($)	0.006	0.016
Flash loan number	1.0	4.0
Flash loan amount ($)	0.016	0.041
Health factor	1.84	1.44
Liquidation number	1.0	2.0
Liquidation amount ($)	0.005	0.033

Table 5: Borrower metrics by margin trading status (median values). All monetary values are reported in millions of USD.

The data further indicate that margin traders operate with both higher frequency and greater volume. On average, they initiate borrowing transactions approximately four times more often and borrow amounts that are three times larger than those of non-margin traders.

In some cases, margin traders may accumulate debt while holding no tokens in their wallets, due to repeated cycles of borrowing and trading. As a result, when the market moves in their favor and they seek to repay debt and reclaim collateral, they are more likely to rely on flash loans to facilitate the transaction. This pattern is reflected in the data, as margin traders use flash loans at a rate four times higher than other users.

While the potential gains from margin trading may be appealing, and flash loans provide a technological mechanism compatible with the objectives of this strategy [Mandin, 2025], the associated risks are considerable. Margin traders tend to maintain lower health factor values and experience liquidation events at twice the rate of non-margin traders.

These findings highlight that margin traders, although a small subset of the Aave user base, play an outsized role. They engage more frequently, take larger positions, and accept greater risk exposure. As a result, they are more prone to liquidation events and exhibit a risk profile that differs markedly from the broader lending community.

5 Liquidations

As discussed above, liquidation is a key risk control mechanism in DeFi lending, which serves as a protective safeguard for lenders. It is designed to prevent lenders from incurring losses when the value of a borrower's collateral deteriorates relative to their outstanding debt. A liquidation is triggered when a borrower's health factor falls below a critical threshold. The health factor is a deterministic function of the value of supplied collateral and the value of borrowed assets. Conceptually, it reflects the distance between the borrower's current position and the point of insolvency. For a formal definition and operational details, see the Aave documentation[11].

Each token listed on Aave is assigned a specific liquidation threshold, which determines the maximum allowable borrowing relative to the collateral value. These parameters are set through protocol governance and are generally more conservative for assets with higher price volatility or lower liquidity. For example, stablecoins are typically assigned higher liquidation thresholds, while more volatile assets such as alternative tokens or wrapped derivatives have lower thresholds. These threshold settings represent a trade-off between protocol safety and efficiency, as overly aggressive liquidation parameters may amplify borrower losses while insufficient thresholds may expose lenders to credit risk [Qin et al., 2021, Cohen et al., 2023].

In the current implementation of Aave V3, a borrower's position becomes eligible for partial liquidation when their health factor falls below one. Under such conditions, up to fifty percent of the debt position may be liquidated. If the health factor deteriorates further and drops below a secondary threshold, currently set at 0.95 by protocol governance (subject to future adjustment), the entire position becomes eligible for full liquidation.

A health factor greater than one indicates that the user remains safe above the liquidation boundary. Users can improve their health factor by either increasing their collateral or partially repaying their borrow position. Since the health factor is directly tied to the market value of the collateral, its sensitivity to price fluctuations introduces significant risk during periods of heightened volatility. A fall in collateral prices leads to a lower health factor and increases the likelihood of liquidation.

The liquidation process is executed by third parties known as liquidators. These actors are incentivized to repay a portion of the borrower's debt in exchange for a discounted portion of the collateral, along with a liquidation bonus [Qin et al., 2021, Cohen et al., 2023]. This decentralized incen-

[11] Accessed on July 29, 2025. Available at: `https://docs.Aave.com/risk/asset-risk/risk-parameters`

tive structure enables undercollateralized positions to be resolved promptly and without the need for centralized enforcement mechanisms. While this design enhances protocol stability under most conditions, it may also introduce unintended consequences [Cohen et al., 2023, Perez et al., 2020].

5.1 Clustering Over Time

While liquidations play a critical role in maintaining system solvency, their frequency and magnitude can signal underlying stress within the lending ecosystem. In an efficient and stable protocol, borrowing and repayment should proceed smoothly, without recurring large-scale disruptions. However, in Aave V3, we observe that liquidation events tend to be highly clustered, occurring in concentrated bursts rather than being evenly distributed over time. Figure 6 illustrates the liquidation amounts over time (red line). The colored areas underneath it show the contributions of the top four tokens.

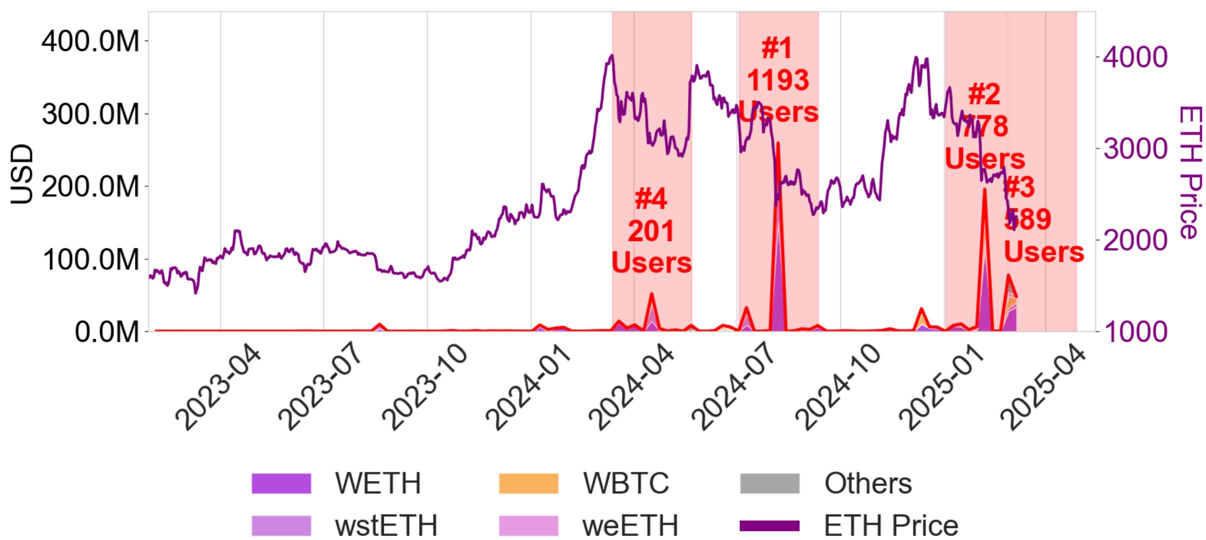

Figure 6: Liquidation amounts over time. The purple line indicate the price of ETH. The red line indicates the total liquidated USD value, while the colored areas underneath it show the contributions of the top four tokens, which together account for approximately 90% of total liquidations.

Importantly, a small number of tokens account for the vast majority of liquidation activity. Specifically, four assets-WETH, wstETH, WBTC, and weETH-jointly account for approximately 90 percent of the total liquidated value in USD. Additionally, we observe that sharp price declines in ETH are strongly associated with spikes in liquidation volume, underscoring the system's sensitivity to major asset price movements.

Outside of these clustered events, liquidations are relatively infrequent and involve modest amounts. This pattern suggests that while liquidation risk is a persistent feature of decentralized lending, it tends to be episodic.

The distribution of liquidations within each wave offers insight into how users with different position sizes are affected. To examine this dynamic, we focus on the ten largest liquidation waves, which together account for roughly 80 percent of total liquidated volume during the sample period. For each wave, we construct a symmetric 48-hour window centered on the peak hour of liquidation activity—capturing the 24 hours before and after the most intense liquidation moment.

To assess the relative impact on users with large versus small positions, we compute, for each hour k, the concentration ratio of liquidated value to the number of individual liquidations within each liquidation wave. This ratio is defined as:

$$\mathrm{VoN}_{t+k} = \frac{\frac{V_{L_{t+k}}}{\sum_k V_{L_{t+k}}}}{\frac{N_{L_{t+k}}}{\sum_k N_{L_{t+k}}}}$$

where $V_{L_{t+k}}$ denotes the liquidated value (in USD), and $N_{L_{t+k}}$ denotes the number of individual liquidation events at hour k within a wave.

This metric serves as a proxy for the concentration of liquidated value. A ratio below one indicates that many small positions are being liquidated, suggesting that retail users or smaller borrowers are most affected. In contrast, a ratio above one implies that fewer, high-value positions are driving the liquidation volume—pointing to greater losses among large borrowers.

As shown in Figure 7, this ratio remains around 0.5 during the twenty-four hour period leading up to the major liquidation event, indicating that users with smaller positions tend to be liquidated earlier. In the two hours immediately preceding the peak, the ratio rises above one, reflecting a shift toward liquidations concentrated in larger positions. Following the peak hour, the ratio drops to near zero, suggesting that liquidation activity declines sharply in both frequency and volume once the primary shock has occurred.

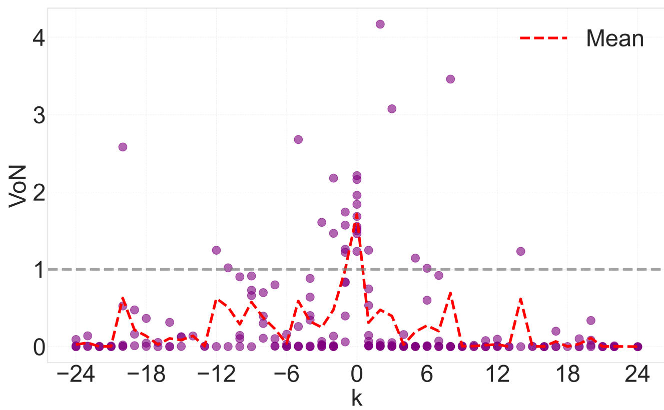

Figure 7: Concentration ratio of liquidated value to the number of individual liquidations for each hour within a 48-hour window (24 hours before and after) within the ten largest liquidation waves. The gray dashed vertical line denotes parity between value and count concentration.

5.2 Cross Sectional Concentration

As highlighted above, liquidations tend to concentrate in distinct waves. The ten largest waves account for approximately 80 percent of total liquidated volume in Aave V3. While the liquidation values during these waves appear large, Table 6 reveals that a disproportionate share of this activity is concentrated among a small number of users.

Most liquidation events, in terms of count, involve users with relatively small positions. However, within each wave, a significant portion of the liquidated value is often attributable to a handful of large borrowers. Columns 3 and 4 of Table 6 provide two measures for the concentration of liquidation. In column 3, the measure CON_1 reports the percentage of the liquidation volume accounted for by the top three users within each wave. For instance, a value of 19 percent indicates that just three users were responsible for 19 percent of the liquidated value during that particular wave. A larger value in this column reflects greater concentration of liquidation among fewer participants.

In Column 4, the measure CON_2 reports the minimum number of users needed to account for 50 percent of the liquidation volume in each wave. A smaller value in this column reflects a greater concentration of losses among fewer participants. Notably, beginning with the sixth wave, the data exhibit a marked increase in concentration, with a few large users driving the majority of liquidated volume.

These results demonstrate that while liquidation volumes may appear substantial in aggregate, they are often concentrated among a small set of borrowers with large positions.

Wave	V_L	N_L	CON_1	CON_2	FEE_L	$OC_{(t+1)}$	$OC_{(t+3)}$	$OC_{(t+5)}$	$OC_{(t+10)}$
2024-08-07	258	2169	19%	15	14 (5%)	14 (5%)	17 (7%)	33 (13%)	38 (15%)
2025-02-06	190	1167	24%	12	11 (6%)	20 (11%)	20 (11%)	20 (11%)	20 (11%)
2025-02-27	76	673	36%	7	5 (7%)	1 (1%)	1 (1%)	3 (4%)	4 (5%)
2024-04-17	48	313	40%	5	3 (6%)	2 (4%)	4 (8%)	4 (8%)	4 (8%)
2025-03-06	47	885	15%	15	3 (6%)	2 (4%)	6 (13%)	6 (13%)	6 (13%)
2024-07-10	32	359	31%	8	2 (6%)	1 (3%)	1 (3%)	2 (6%)	4 (13%)
2024-12-12	30	88	81%	2	2 (7%)	1 (3%)	3 (10%)	3 (10%)	4 (13%)
2024-03-19	14	82	71%	2	1 (7%)	0 (0%)	1 (7%)	1 (7%)	1 (7%)
2025-01-16	10	165	57%	3	1 (10%)	1 (10%)	1 (10%)	1 (10%)	2 (20%)
2023-08-20	9	74	62%	2	1 (11%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)

Table 6: Summary of liquidation waves and borrower Information on Aave V3. V_L denotes the total liquidated value, and N_L represents the number of individual liquidation events. CON_1 and CON_2 are the two measures of liquidation concentration. FEE_L refers to the liquidation fee. OC refers to the estimated opportunity cost. All monetary values are reported in millions of USD. Percentages reported for fees and opportunity cost are expressed relative to $L(\$)$

5.3 Liquidation Triggers

A position becomes eligible for liquidation when the user's health factor falls below one. In this section, we decompose the effects that cause the health factor to drop to that level. The health factor (HF) on Aave is defined as:

$$\text{HF} = \frac{\text{Total Collateral in USD} \times \text{Portfolio Liquidation Threshold}}{\text{Total Debt in USD}} = \frac{\left(\sum_{i\in S} p^i q^i r^i\right)\bar{LT}}{\sum_{i\in S} p^i Q^i R^i},$$

where p^i denotes the market price of token i, q^i and Q^i denote the scaled quantities of token i held as collateral and debt, respectively, and r^i and R^i denote the gross accrued supply and borrow rates. The set S contains all tokens in the user's portfolio. We let $\bar{LT}$ denote the effective portfolio liquidation threshold, defined as the collateral-value-weighted average liquidation threshold across the tokens in the user's collateral portfolio. Since token-specific liquidation thresholds are protocol-defined parameters that change only through governance and such changes are extremely rare over our sample period, we treat $\bar{LT}$ as approximately constant over the event window considered in the decomposition.

Under this approximation, changes in the health factor are driven primarily by changes in token prices, quantities, and interest rates. In this section, we decompose the change in health factor into three components:

1. changes in token prices impacting both the values of collateral and debt
2. changes in interest rates impacting both the values of collateral and debt
3. changes in token quantities due to borrowing and supplying (B&S) activities, including repaying and withdrawing.

Additionally, we further disaggregate the price effect into components driven by collateral price movements and debt price movements.

It is important to emphasize that our decomposition identifies mechanical contributors to health factor deterioration but does not attribute causality to the underlying market behavior or trading motives.

Assuming that the liquidation threshold (LT) remains fixed, we derive the decomposition first by calculating the collateral and debt weights:

$$w_C^i = \frac{p^i \cdot q^i \cdot r^i}{\text{Total Collateral}}, \quad w_D^i = \frac{p^i \cdot Q^i \cdot R^i}{\text{Total Debt}}.$$

First-order approximation of the relative change in health factor then can be written as:

$$\frac{dHF}{HF} \approx \underbrace{\sum_{i \in S} (w_C^i - w_D^i) \cdot \frac{dp^i}{p^i}}_{\text{Price Effect}} + \underbrace{\sum_{i \in S} w_C^i \cdot \frac{dq^i}{q^i} - \sum_{i \in S} w_D^i \cdot \frac{dQ^i}{Q^i}}_{\text{Quantity Effect}} + \underbrace{\sum_{i \in S} w_C^i \cdot \frac{dr^i}{r^i} - \sum_{i \in S} w_D^i \cdot \frac{dR^i}{R^i}}_{\text{Rate Effect}}.$$

Hence, the decomposition of the price effect is derived as follows:

$$\text{Price Effect} = \underbrace{\sum_{i \in S_C} (w_C^i - w_D^i) \frac{dp^i}{p^i} + \sum_{i \in S_{CD}} (w_C^i - w_D^i) \frac{dp^i}{p^i} \cdot \frac{1}{2}}_{\text{Collateral Price Effect}} + \underbrace{\sum_{i \in S_D} (w_C^i - w_D^i) \frac{dp^i}{p^i} + \sum_{i \in S_{CD}} (w_C^i - w_D^i) \frac{dp^i}{p^i} \cdot \frac{1}{2}}_{\text{Debt Price Effect}}.$$

Here, the sets S_C, S_D, and S_{CD} denote the tokens used exclusively as collateral, exclusively as debt, and jointly as both collateral and debt, respectively. These sets are mutually disjoint, and their union gives the full set of tokens (S) in the user's portfolio. This decomposition allows for a precise attribution of health factor changes. While this analysis does not explore the structural drivers behind token price or rate changes, it provides a transparent framework to assess the proximate causes of liquidation triggers.

The above decomposition framework helps estimate the drivers of health factor deterioration immediately before a liquidation. Specifically, we examine the changes in collateral prices, debt prices, interest rates, and net token movements during the one-hour period preceding each user's largest liquidation event. This analysis is conducted for two groups: (1) the top ten users with the largest individual liquidations, and (2) all users who experienced liquidation events. The results are summarized in Table 7.

Group\Triggers	Collateral price	Debt price	Interest rate	B&S activities
Top 10 Users	97.28%	2.44%	0.29%	0.00%
All Users	84.40%	10.59%	4.05%	0.60%

Table 7: Weighted average of proportional effects on the health factor one hour prior to each user's largest liquidation event. Values are weighted by user debt prior to liquidation.

For both groups, the predominant driver of liquidation is the change in the market value of collateral. Among the top ten users, this effect accounts for approximately 97 percent of the observed deterioration in the health factor. For the broader set of liquidated users, the figure is slightly lower at 84 percent. The impact of debt price movements is more modest, contributing around 2 percent for the top ten users and 11 percent for the overall group.

The contribution of accrued interest rates to health factor decline is minimal in both samples, suggesting that changes in protocol-level rates play only a secondary role in immediate liquidation risk. Likewise, trading activity-including borrowing, repaying, supplying, and withdrawing-has negligible influence on health factor changes in the critical hour prior to liquidation.

Taken together, these findings suggest that liquidation events are primarily triggered by rapid declines in the value of collateral, rather than increases in the value of debt or user-initiated portfolio adjustments. The typical liquidation scenario therefore involves users who hold volatile collateral assets and are vulnerable to drastic changes in price of the collateral, with only a minor portion of the liquidation pressure attributable to rising debt valuations.

5.4 Estimation of Borrowers' Loss

According to Table 6, the total liquidated value is over 250 millions for the largest liquidation wave. While liquidation volumes can appear large during major waves, the actual financial losses incurred by borrowers are often less clearly understood. In this section, we outline our methodology for estimating borrower losses resulting from liquidation events.

To accurately interpret the mechanics of liquidation, a conceptual clarification is necessary. Liquidation is not inherently a loss-generating event in the conventional sense; rather, it is a forced settlement process in which a borrower's collateral is sold to repay their outstanding debt at a specific point in time. The liquidated value reflects the portion of debt repaid through the sale of collateral, but it does not directly quantify borrower losses. Instead, borrower loss should be understood as the gap between what is forcibly repaid and what the borrower might have retained under more favorable market conditions—such as if they had avoided liquidation and benefited from a subsequent price recovery.

We estimate borrower loss through two distinct channels:

1. **Liquidation Fee (Fee_L)**: When a borrower's collateral is liquidated, a liquidation penalty, typically in the form of a fee, is added to the amount of collateral seized. This fee is not returned to the borrower and thus constitutes an immediate and direct financial loss.

2. **Opportunity Cost (OC)**: In addition to the fee, borrowers may incur a loss from being forced to sell collateral at a suboptimal price. If the price of the liquidated asset subsequently increases, the borrower misses the opportunity to benefit from the price recovery. We define this as the opportunity cost.

To estimate opportunity cost, we proceed as follows. Let p_t^i denote the price of collateral token i at the time of liquidation t, and let q_t^i represent the liquidated quantity. We then define the opportunity cost over an Δ day window as:

$$\mathrm{OC}^i_{(t+\Delta)} = \max_{\tau \in [t, t+\Delta]} (p^i_\tau - p^i_t) q^i_t,$$

where p_τ^i denotes the price of token i at any time τ between the liquidation time t and the end of the Δ day window. This metric captures the maximum unrealized gain a borrower could have obtained had they retained the collateral through the price recovery.

The last five columns of Table 6 report the results of this estimation. Liquidation fees range between 5 and 10 percent of the liquidated value across waves. In contrast, estimated opportunity cost varies more widely. Within a one-day window following liquidation, the opportunity cost ranges from near zero to approximately 10 percent in some waves. When the observation window is extended to three days, opportunity cost rises above 10 percent in several cases and remains relatively stable when the window is further extended to five or ten days.

These findings suggest that asset prices often recover following liquidation events. As a result, borrowers may suffer combined losses, including both liquidation fees and foregone gains, on the order of 10 to 30 percent of the liquidated value, relative to the best-case counterfactual scenario in which they had retained their collateral through the recovery.

5.5 Impacts on Market Prices

Liquidation mechanisms are intended to protect lenders by ensuring that undercollateralized positions are promptly resolved. However, these mechanisms may also generate unintended consequences for broader market dynamics. In particular, clustered liquidation events could exert persistent pressure on market prices, generating a "fire-sale" externality to other market participants.

This section presents empirical evidence to evaluate the claim and whether liquidations on Aave V3 exert a persistent impact on market prices. To do so, we account for several confounding and scale-related factors. If a persistent price effect is observed, it may be due to external factors correlated with the liquidation event rather than the liquidation itself. To address this, we introduce a control group, defined as tokens that were not subject to liquidation. Additionally, to capture the relative market significance of a liquidation, we normalize liquidation amounts by the respective token's market capitalization. Finally, since both wave-specific and token-specific effects may influence price behaviour, we introduce corresponding fixed effects to disentangle these influences from the liquidation shock itself.

Given these concerns, we estimate the following regression model to analyze how liquidation activity relates to asset price movements:

$$\frac{p^i_{t+k} - p^i_{t-1}}{p^i_{t-1}} - \frac{p^{\text{Control}}_{t+k} - p^{\text{Control}}_{t-1}}{p^{\text{Control}}_{t-1}} = \alpha + \beta \cdot \frac{L^i_t}{MC^i_t} + \sum_{t=1}^{9} \gamma_t \cdot D_t + \sum_{i=1}^{M} \delta^i \cdot D^i + \epsilon,$$

where p^i_{t+k} is the price of token i, k days after the liquidation event at time t, L^i_t is the liquidation volume of token i during wave t, and MC^i_t is the market capitalization of token i at time t. The left-hand side of the equation captures the relative price change of the liquidated token net of a control-group benchmark. Dummy variables D_t and D_i control for time-specific and token-specific fixed effects, respectively.

To isolate the effect of liquidation from broader market movements, we employ two control-group benchmarks: (1) Litecoin (LTC), which is not used as a collateral in DeFi lending and hence not subject to liquidations, and (2) CCI30, a market-capitalization-weighted index tracking the top thirty cryptoassets.

The analysis is limited to the top ten liquidation waves, which jointly account for approximately 80 percent of total liquidation volume. While this restricts the sample, it provides sufficient coverage to investigate price effects.

The hypothesis underlying this design is that any observed price pressure following liquidation may be attributable to a common factor affecting both the liquidated tokens and broader market conditions, rather than the liquidation itself. We estimate the model for values of k ranging from 0 to 20, capturing price changes from the day of liquidation up to twenty days after.

Figure 8 presents the estimated β coefficients and 95 percent confidence intervals across three specifications: (1) no benchmark adjustment, (2) LTC as benchmark, and (3) CCI30 as benchmark. Across all specifications, we find no statistically significant evidence that the relative size of liquidation events has a persistent effect on market prices.

To complement this analysis, we also examine the estimated intercepts from the same models, which capture unconditional average price changes following liquidation events. Figure 9 reports the intercepts and corresponding 95 percent confidence intervals. In the specification without any benchmark adjustment, we observe a short-term negative price pressure in the first four to six days following a liquidation event. However, this effect becomes statistically insignificant once we control for either LTC or the CCI30 index, suggesting that any observed price declines may reflect general market conditions rather than liquidation-specific effects.

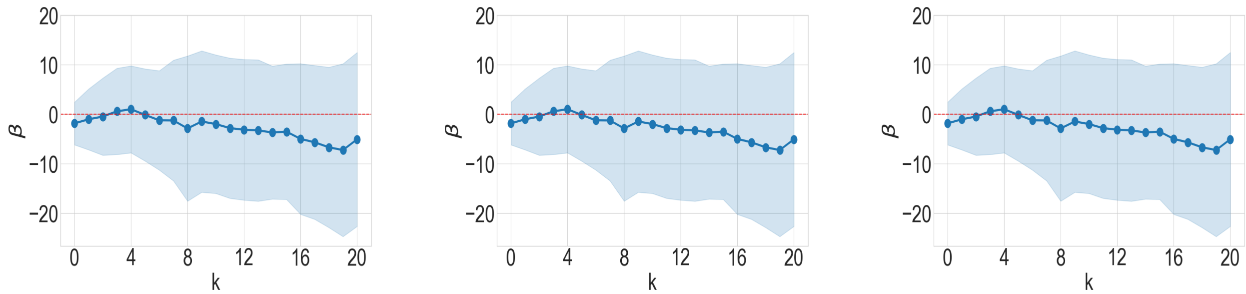

No control | Control = LTC | Control = CCI30

Figure 8: Estimated β coefficients and 95 percent confidence intervals for relative liquidation size across days $k = 0$ to $k = 20$, under three benchmarks.

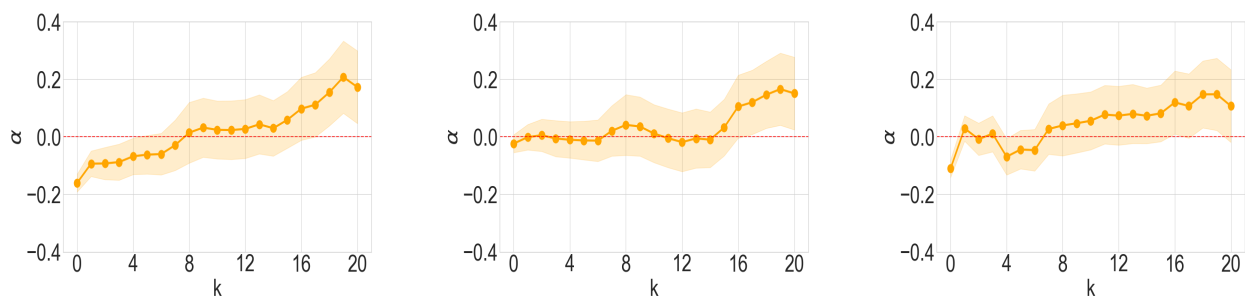

No control | Control = LTC | Control = CCI30

Figure 9: Estimated intercepts and 95 percent confidence intervals for price effects across days $k = 0$ to $k = 20$, under three benchmark.

Taken together, the results provide no statistically significant evidence that liquidation activity leads to a persistent effects on market prices. While short-term impacts may be present, these effects are not robust to benchmark controls. This suggests that the observed price dynamics following liquidation may be driven by broader market forces that simultaneously affect both liquidated tokens and benchmark assets. Nevertheless, we acknowledge the limitations of the presented model, particularly its limitation to capture more nuanced or indirect effects of liquidation on market prices. Further research is necessary to validate these findings and provide more robust evidence on the potential market impact of liquidation events.

6 Conclusion

This paper examines the mechanics and implications of decentralized lending using granular data from Aave V3, the largest DeFi lending protocol by total value locked. We document how the platform generates revenue, identify margin trading behavior among borrowers despite overcollateralization requirements, and analyze the triggers, dynamics and consequences of liquidation events. Together, these findings provide insight into how DeFi lending protocols function in practice and how they differ from traditional financial intermediation.

Our analysis suggests that lending without traditional intermediaries is viable in a technical and operational sense. Aave V3 successfully matches borrowers and lenders, enforces collateral constraints through smart contracts, and maintains solvency without relying on trust, identity, or centralized enforcement. The system operates continuously, transparently, and with minimal overhead, demonstrating the potential of protocol-based credit markets.

At the same time, important limitations remain. First, we observe persistent underutilization of deposited funds, reflecting structural inefficiencies and fragmentation across risk pools. Second, overcollateralization requirements impose a high capital cost on borrowers and limit its usefulness for real-world economic activities. Third, the reliance on liquidation as a primary risk control mechanism can result in sudden and significant losses, especially during periods of market volatility. Finally, the widespread use of recursive leverage—enabled by frictionless rehypothecation of collateral—can amplify market movements and introduce systemic fragility within the crypto ecosystem, even in the absence of direct interconnectedness among users.

These limitations suggest a potential role for the public sector in shaping the future evolution of decentralized lending. One direction involves the tokenization of real-world assets (RWAs), which could expand the collateral base and improve the quality and stability of lending markets. Another is the development of decentralized identity frameworks that allow for the selective disclosure of user attributes, enabling better underwriting without fully compromising anonymity. More broadly, it would be interesting at least academically to explore the desirability and feasibility of prudential regulations—including capital requirements, leverage limits, or liquidity thresholds—to DeFi protocols. Could these regulations help align platform incentives with systemic stability while preserving the benefits of on-chain automation? Of course, important open questions remain about the feasibility of designing and enforcing such rules given pseudonymous participation, cross-border activity, and limited direct control over smart contracts.

References

[Chiu et al., 2023] Chiu, J., Ozdenoren, E., Yuan, K., and Zhang, S. (2023). On the fragility of defi lending. Staff Working Paper 2023-14, Bank of Canada. 1

[Cohen et al., 2023] Cohen, S. N., Sabate-Vidales, M., Szpruch, L., and Gontier Delaunay, M. (2023). The paradox of adversarial liquidation in decentralised lending. *SSRN Electronic Journal*. 5

[Heimbach and Huang, 2024] Heimbach, L. and Huang, W. (2024). Defi leverage. BIS Working Paper 1171, Bank for International Settlements. 1

[Heimbach et al., 2023a] Heimbach, L., Schertenleib, E. G., and Wattenhofer, R. (2023a). Short squeeze in defi lending market: Decentralization in jeopardy? 9

[Heimbach et al., 2023b] Heimbach, L., Seiler, V., and Spann, M. (2023b). Defi lending: Intermediation without intermediaries? CESifo Working Paper Series 10045, CESifo. 1

[Kozhan et al., 2024] Kozhan, R., Liu, D., Mueller, M. M., and Zhong, M. (2024). Defi lending and leverage. Working paper. 1

[Lehar and Parlour, 2022] Lehar, A. and Parlour, C. A. (2022). Systemic fragility in decentralized markets. BIS Working Paper 1062, Bank for International Settlements. 1

[Mandin, 2025] Mandin, J. (2025). Risk-free uncollateralized lending in decentralized markets: An introduction to flash loans. Staff Discussion Paper 2025-6, Bank of Canada. 4

[Park and Stinner, 2023] Park, A. and Stinner, J. (2023). Phantom liquidity in decentralized lending. *SSRN Electronic Journal*. 4

[Perez et al., 2020] Perez, D., Werner, S. M., Xu, J., and Livshits, B. (2020). Liquidations: Defi on a knife-edge. 1, 5

[Qin et al., 2021] Qin, K., Zhou, L., and Gervais, A. (2021). An empirical study of defi liquidations: Incentives, risks, and instabilities. 1, 5

[Rivera et al., 2023] Rivera, T., Saleh, F., and Vandeweyer, Q. (2023). Equilibrium in a defi lending market. Available at AFA 2024. 1

[Warmuz et al., 2023] Warmuz, J., Chaudhary, A., and Pinna, D. (2023). Toxic liquidation spirals. 9