

The Structural Shift in DeFi Lending

Ole E. Pedersen^{1,2}, Johannes Rude Jensen^{1,2}, and Omri Ross^{1,2}

¹ Department of Computer Science, University of Copenhagen, Copenhagen, Denmark

² eToro Labs, Copenhagen, Denmark
olpe@di.ku.dk, johannesje@etoro.com, omri@di.ku.dk

Abstract. DeFi lending markets have grown exponentially in recent years. We document an increasing level of sophistication among DeFi lenders and borrowers, demonstrating a structural shift from simple directional bets toward highly leveraged, yield-driven strategies, predominantly in correlated collateral-debt pairs. Using on-chain data for Aave v3 and SparkLend on Ethereum from January 2023 to December 2025, we reconstruct address-level collateral, debt, and implied leverage across 59 assets, capturing roughly two-thirds of aggregate TVL for the period. We infer strategies from collateral-debt compositions and demonstrate how recursive lending with staked ETH and yield-bearing stablecoins has become the dominant strategy. At the peak, in August 2025, over 65% of outstanding debt was tied to correlated positions backed by less than 15% of aggregate equity. These strategies materially reshape interest rate dynamics of on-chain lending markets, while concentrating credit risk on assumptions of withdrawal liquidity and correlation. We discuss the implications of these findings and provide directions for future research.

Keywords: Decentralized finance · DeFi lending · Leverage · Liquid staking · Recursive lending

1 Introduction

Decentralized finance (DeFi) lending protocols such as Aave, Compound, Morpho, Sky, and Euler are systems composed of interoperable smart contracts deployed on public blockchains, predominantly on Ethereum and its associated L2s. Participants interact with smart contracts directly through externally owned accounts (EOAs) or via other decentralized applications (dApps), enabling permissionless liquidity provision and collateralized credit access.

These systems operate as two-sided markets between liquidity suppliers and borrowers. Suppliers deposit tokens or assets, receiving interest. Borrowers post overcollateralized positions, drawing liquidity against their pledged assets. The absence of identity-based credit assessment implies that credit risk for the protocol must be managed through collateral, and the liquidation of said collateral must be automated to mitigate bad debt when sudden spikes in volatility reprice collateral against debt. Hence, each available asset is parameterized by a set of specific risk parameters, such as the collateralization ratio and liquidation

threshold. The fundamental relationship between supply and demand for an asset determines its utilization rate, which dynamically drives both borrowing and lending yields.

In equilibrium, the interest rate functions for the protocol align to maintain near-optimal capital efficiency without exhausting available liquidity. Accordingly, to maintain solvency and mitigate credit risk, borrowing positions must remain overcollateralized. Each asset is assigned a loan-to-value (LTV) limit and a liquidation threshold that determines when automated liquidations occur. The ratio of collateral value to debt defines the solvency indicator (“health factor”) for a given market. When market volatility or price shocks push a loan into insolvency, arbitrageurs are economically incentivized to repay the borrower’s debt in exchange for discounted collateral. As the vast majority of liquidations are purchased by MEV bots attempting to collect the maximal extractable value (MEV) for each successive block, liquidations are typically processed within the span of a single block [13].

Historically, lending markets have served two purposes in DeFi: first, enabling committed long-term holders to access capital without losing exposure by facilitating lending of stablecoins against crypto. Second, borrowing for leveraged exposure by recursively borrowing and re-collateralizing assets, a practice colloquially referred to as *looping*. When looping assets, an agent repeatedly performs a sequence of operations in which existing collateral is used to obtain additional borrowing capacity, and the proceeds of that borrowing are reinvested as collateral. So long as the agent borrows up until the limit determined for the collateral, the sequence constitutes a geometric progression with a common ratio.

The cumulative debt is the sum of all intermediate borrowings, while the total exposure is the sum of the collateral locked in the lending protocol. In other words, looping entails going long the collateral while shorting the borrowed asset. In most cases, the agent will choose to borrow a stablecoin in order to ensure a stable principal against the (assumed) appreciating value of the collateral. Hence, as the number of loops increases, the position becomes progressively more sensitive to adverse price movements, since both collateral and debt expand geometrically while the solvency margin remains fixed. Until recently, the natural limitation to recursive lending was provided by the liquidation threshold given by the assumed safe loan-to-value (LTV) margin for an asset. While infinite recursion is theoretically possible, the maximum exposure achievable is asymptotically bounded by the LTV ratio.

Essentially, the size at which another recursion provides meaningful results quickly converges, especially when gas costs are factored in. To see this, consider an asset with an LTV of 80%. Here, the practical limit to additional recursion would be capped below $5 \times$ notional exposure. However, with the recent addition of restricted collateral regimes in lending markets (e.g., *eMode* on Aave or correlated vaults on Morpho), protocols enable the application of a distinct set of risk parameters to assets exhibiting high historical correlation, under the assumption that intra-category price volatility remains low. This feature will typically only be enabled for tightly correlated assets, such as a base asset (ETH/WETH)

and derivative tokens compounding staking yield over time (wstETH) (see Appendix B). Theoretically, this promotes capital efficiency for correlated asset classes while constraining exposure to prevent systemic spillover, by partitioning borrowing capacity into quasi-isolated risk silos. In practice, the increased liquidation thresholds (typically upwards of 95%) for these tightly correlated asset classes enable recursive leveraging towards $20\times$ notional exposure.

In this paper, we demonstrate how users of on-chain lending markets are increasingly employing leveraged strategies, relying on the increased liquidation thresholds to borrow recursively. We show how loans, where the collateral and underlying token exhibit historically tight correlation, are used to maximize exposure to yield-bearing derivative tokens or stablecoins, enabling agents to collect the spread between the cost of the loan and the yield. We note that this paper makes use of technical terminology related to the specific design and implementation of each token contract. For an introduction and overview of each asset category and the underlying economic model, we invite the reader to review Appendix A prior to reading the paper.

2 Literature

Recent research has increasingly focused on how decentralized lending protocols behave when agents pursue leveraged strategies of recursive borrowing with correlated collateral. While early empirical results showed that DeFi borrowing was primarily motivated by directional exposure to ETH- or BTC-variants [15,10], subsequent studies documented a gradually emerging tendency for yield-arbitrage, initially to capture yield-farming incentives [17]. It was only later, after “The Merge,” with the emergence of liquid-staking derivative tokens, that recursive leverage took on the relative proportions documented in this paper [2].

Recent years have seen an increasing interest in how recursive leverage in correlated assets produces geometric leverage that captures interest rate spreads [15], under the assumption that the correlations will hold [2,13,16]. This practice, combined with the rehypothecation of restaking, inflates the “total value locked” (TVL) figures frequently used in the industry, which can obscure how the concentration of credit risk on oracle or smart-contract infrastructure compounds with the size of the notional exposure on lending markets [11].

As noted above, the introduction of Aave’s eMode and Morpho’s correlated vaults has further amplified this tendency. By assuming stable intra-category price relationships, these markets enable significantly higher LTV ratios than previously allowed [4] [16]. While higher LTVs for historically correlated assets can increase capital efficiency for the market, the adverse implication is triggering a rapid deleveraging sequence of cascading liquidations, if correlations weaken or interest rates become unfavorable.

Although staking derivatives generally remain tightly correlated with their underlying token, recent work has underscored how staking derivatives, such as wstETH (Appendix A), can deviate from their underlying token due to multiple endogenous reasons, including liquidity constraints, funding frictions, and

withdrawal queues, especially under stress [14,16]. Hence, when staking derivatives make up the lion’s share of collateral in lending protocols, even small peg deviations can propagate nonlinearly and trigger widespread liquidations [2,16].

For this reason, the liquidation mechanisms and the ‘liquidation pipeline’ (post-liquidation) have been the target of scrutiny. The automated liquidation-auctions in lending protocols are currently dominated by MEV bots, given the direct incentive for capturing the arbitrage between the liquidation penalty and the price in AMMs such as Curve or Uniswap. This has been shown to lead to timing asymmetries in which debt is repaid instantly while collateral is seized at a discount [13]. Due to the underlying logic of lending pools, edge-case vulnerabilities around oracle updates can be manipulated to trigger forced liquidations [3]. This risk is especially pertinent for highly leveraged staking derivatives, as small price shocks can liquidate large portions of collateral, often within a single block [12].

3 Methodology

Aave v3 and SparkLend (forked from Aave v3) issue ERC-20 tokens to represent both collateral supplied and debt incurred. When a user deposits assets, whether to lend or to post as collateral, the protocol mints a derivative token that encodes ownership of the underlying funds and can be redeemed only by returning the token. Similarly, debt positions are represented by distinct ERC-20 tokens that remain in the borrower’s possession until the loan and its associated interest are repaid. Each token type uniquely corresponds to its underlying asset. By reconstructing all historical transfers of these tokens to and from each address, we infer daily balances of collateral and debt across all assets.

3.1 Data Collection and Filtering

The dataset contains the full set of event logs for all interactions with the two lending protocols. The ratio between logs heavily favors Aave v3, as this market remains the largest by far.

Table 1. Summary statistics for Aave v3 and SparkLend on Ethereum after filtering.

Aave v3 SparkLend		
Tokens	58	18
Unique borrowers	61,084	5,768
Days	1,39	978
Borrow transactions	584,913	28,599
Liquidations	17,883	1,769

We collect blockchain transactions for all addresses interacting with Aave v3 and SparkLend between January 2023 and December 2025, using data obtained

from Dune Analytics. The dataset contains 59 token contracts (Appendix B), and on average for the period, it captures 65.1% of the total value locked (TVL) for all on-chain lending protocols on Ethereum throughout the observation period [8].

Aave and SparkLend maintain oracle smart contracts that return a token’s current price when provided with its address. Because these contracts do not emit price-update events, our event-log dataset lacks historical price information. To collect pricing data, we query an archive node via Infura, allowing interactions with smart contracts at any past state. By querying the archive node at daily intervals over the observation period, we construct a time series of token prices in U.S. dollars, reflecting the values used by Aave and SparkLend.

We make the simplifying assumption that all users on Aave borrow at the variable borrow rate. While Aave previously also allowed users to borrow at a stable rate, this assumption enables us to drastically reduce the complexity when determining accrued interest. Aave disabled the stable borrow rate in late 2023, so this assumption affects less than 2.6% of the total debt examined in this analysis. Consequently, even if all debt before this date was borrowed at the stable rate, the effect on total debt would be well below $\pm 0.1\%$, hence we assume that this has negligible impact on our results.

Finally, we filter the data as follows. First, the address must be solvent; any address whose outstanding debt exceeds its collateral is excluded. Second, the address must maintain debt amounting to at least five percent of its collateral value, thereby removing addresses whose activity more closely resembles passive lending rather than strategic or recursive borrowing. Third, users must have incurred more than one U.S. dollar in debt, ensuring that trivial “dust” balances and minor rounding discrepancies do not contaminate the analysis.

3.2 Data Analysis and Definitions

For Aave and SparkLend, a naïve token-balance measure would fail to capture interest accrual. To address this, we normalize all historical transfers by the asset-specific liquidity index at the time of transfer. The liquidity index is a monotonically increasing variable that tracks cumulative interest. Multiplying the normalized balance by the current liquidity index yields the present value of a user’s position, inclusive of accrued interest. Equations (1a) and (1b) demonstrate the mechanism used for determining a user’s balance of collateral or debt tokens at any given time t .

$$C_t = \sum_{i < t} \frac{S_i}{L_{C_i}} \times L_{C_t} \quad (1a) \quad D_t = \sum_{i < t} \frac{B_i}{L_{D_i}} \times L_{D_t} \quad (1b)$$

The number of tokens posted as collateral by a borrower, C_t , is determined from a sequence of deposits and withdrawals denoted by S_i , with positive values indicating deposits and negative values indicating withdrawals. Likewise, the number of debt tokens, D_t , is defined in terms of borrowings and repayments

denoted by B_i . Finally, L_i represents the liquidity index at the time when each flow occurs, and L_t when we perform the daily snapshot.

Having reconstructed all collateral and debt balances throughout the observation period, we compute each remaining address's implied leverage. Any borrowed funds inherently introduce leverage, regardless of whether the borrowed assets are spent externally or reinvested into additional cryptocurrency exposure. As described above, a protocol's collateralization requirements impose an upper limit on feasible leverage, determined by the loan-to-value (LTV) ratio, as expressed in Eq. (3).

$$\text{Maximum Leverage}_t = \frac{1}{1 - LTV} \quad (3)$$

Historically, achieving leverage near this theoretical limit required manually swapping borrowed assets for collateral and re-borrowing (looping). Innovations such as non-custodial automated risk management tooling or even flash loans now enable users to automate the management of leveraged positions, and to establish highly leveraged positions more efficiently. Nevertheless, in practice, we observe that most users still maintain leverage well below the protocol's maximum. This realization is achieved by applying the standard equity-multiplier in Eq. (4) to quantify actual leverage.

$$\text{Leverage}_t = \frac{C_t \times P_{C_t}}{C_t \times P_{C_t} - D_t \times P_{D_t}} \quad (4)$$

In Eq. (4), $C_t \times P_{C_t}$ and $D_t \times P_{D_t}$ represent the notional value of a single collateral-debt pair, both measured in U.S. dollars at time t . This measure intentionally abstracts away directional market exposure. If collateral and debt are denominated in the same asset, or in assets that are historically highly correlated, the user's net delta exposure may be minimal even when the measured leverage is high. Conversely, oppositely correlated positions may produce significantly greater delta exposure than the formula implies. This is particularly salient for asset pairs such as wstETH and WETH, which share the same underlying asset (ETH) yet diverge over time due to the embedded staking yield in wstETH. Despite these limitations, leverage remains a crucial indicator for characterizing borrower behavior.

4 Findings

First, we examine the aggregate collateral, debt, and equity across the two lending protocols. We start by grouping all users into one of four asset categories, each assigned based on the assets they borrow. The three main categories are ETH-variants, USD-variants, and BTC-variants. A user is assigned to one such group if the asset variants constitute at least 90% of their total debt. Finally, the "Others"-variants contain all other users not suited for any category.




Fig. 1. Aggregate collateral, debt, and equity for borrowers for whom the asset class makes up more than 90% of their debt. The groupings are **ETH**: borrowers with 90% of their debt in **ETH**-variants; **USD**: borrowers with 90% of their debt in stablecoins; **BTC**: borrowers with 90% of their debt **BTC**-variants; **Others**: all other borrowers that do not fall under the other categories. Note that the majority of debt on lending markets is **ETH**- or **USD**-variants, as the vast majority of assets fall within these categories.

As is evident from Figure 1, both collateral and debt are dominated by ETH- and USD-variants. Notably, at the peak, USD-variant borrowers account for roughly 50% of the debt but around 90% of equity, while ETH-variant borrowers account for the remaining roughly 50% of the debt but only about 10% of the equity. By looking more closely at the split between debt and collateral for the individual assets, we find a plausible explanation of this tendency (Figure 2). Evidently, WETH borrowers predominantly engage in one of two strategies: (i) borrow WETH and buy staked or restaked ETH, or (ii) borrow staked ETH and buy restaked ETH.




Fig. 2. Collateral and debt distribution for two mutually exclusive subsets of users borrowing ETH-variants. The users have been assigned to $\mathbf{ETH} \rightarrow \mathbf{ETH}_{\text{LST}}$ if 50% of their debt is unstaked ETH, i.e. WETH. If not, they are assigned to $\mathbf{ETH}_{\text{LST}} \rightarrow \mathbf{ETH}_{\text{LST}}$. Note how this heuristic reveals a clear preference for borrowing ETH-variants of lower yield to buy ones of higher yield. Assets with less than 1% of the total debt or collateral have been removed from this figure for visualization purposes.

Both strategies express a clear preference for maximizing the staking yield for ETH, while only having low collateral exposure. The key assumption here is the sustained correlation between WETH and the underlying ETH locked in its yield-bearing variants, primarily liquid-staking derivatives (weETH, wstETH, etc.). For the USD-variants, a similar tendency emerges. Presumably, stablecoin borrowers predominantly engage in one of two strategies: (i) borrow non-yield-bearing stablecoins to buy yield-bearing stablecoins, or (ii) borrow non-yield-bearing stablecoins and buy ETH- or BTC-variants. The former attempts to maximize the yield from stablecoins, while the latter expresses a preference for increasing directional exposure to the collateral token.

Extending this preference, we note the considerable popularity of USD-variant PT tokens from the Pendle platform. PT tokens are the principal leg of a yield-bearing asset, stripped of all future yield and redeemable at expiry for the underlying (without the yield). Essentially, these tokens behave like zero-coupon instruments whose discounted price deterministically converges to par. By looping PT tokens with USD-variants (stablecoins) at a lower floating rate, the agent's

payoff is the spread between the PT's implied fixed yield and the interest paid on the loan.

Lastly, when viewing non-USD-variant assets collateralized by addresses borrowing USD-variants (stablecoins), we see the well-documented strategy of leveraging BTC- or ETH-variants (both yield-bearing and non-yield-bearing) primarily against non-yield-bearing USD-variants (stablecoins).




Fig. 3. Collateral and debt distribution for two mutually exclusive subsets of users borrowing USD-variants. The users have been assigned to **USD → USD** if 90% of their collateral is in any kind of stablecoin, potentially yield-bearing. If not, they are assigned to **USD → Ex-USD**. Note how this heuristic shows a clear tendency for either borrowing non-yield-bearing stablecoins (USDT/USDC) to buy yield-bearing USD-variants, or for borrowing non-yield-bearing stablecoins (USDT/USDC) to buy ETH- or BTC-variants. Assets with less than 1% of the total debt or collateral have been removed from this figure for visualization purposes.

To understand how the tendency for these sophisticated strategies has developed over time, we differentiate between the most prevalent strategies in the dataset (Figure 4). To do so, we construct six mutually exclusive categories, collectively covering all the assets in the dataset.

Note the dominance of correlated debt–collateral strategies. Despite only accounting for 15% of the equity, they consume over 65% of the total debt at their peak, prior to the market-wide deleveraging event in October 2025. Before this point, debt levels of these strategies reached more than \$20 bn, but they appear to have declined considerably since then. However, if we measure the quantity of ETH derivatives being borrowed, we find that the debt only declined by 20.9% from its highest point on July 15 to its lowest point on November 20, although the USD-denominated debt shown in Figure 4 suffered a 52.2% decline over the same period. Therefore, while it appears the amount of debt used for extracting yield from tightly correlated tokens has declined considerably since its peak, we ascribe the weakening in the price of ETH relative to the U.S. dollar to be the leading cause, rather than drastically lower utilization of this strategy.




Fig. 4. Aggregate debt and equity across six mutually exclusive categories. The categories are **ETH_{LST} → ETH_{LST}**: strategies that borrow yield-bearing ETH-variants to collateralize yield-bearing ETH-variants; **ETH → ETH_{LST}**: strategies that borrow non-yield-bearing ETH-variants to collateralize yield-bearing ETH-variants; **USD → USD**: strategies that borrow USD-variants to collateralize USD-variants; **USD → Ex-USD**: strategies that borrow USD-variants to collateralize ETH, BTC or "Others"-variant tokens; **BTC**: strategies that borrow BTC-variants to collateralize any token; **"Others"**: strategies that borrow any remaining asset types to collateralize positions not covered by the categories above. Note that borrowing tokens to engage in arbitrage of correlated debt and collateral has become increasingly dominant and accounts for upwards of 65% of the total debt. At its peak, this debt is only secured by less than 15% of the equity.

Next, we examine the implied leverage multiples of the six dominant strategies defined previously. We constructed the leverage distribution by aggregating the debt of all users sharing the same leverage multiples, then repeated this for each daily snapshot over the entire observation period. Figure 5 clearly demonstrates how the strategies involving tightly correlated pairs apply considerably higher leverage compared to the directional strategies.




Fig. 5. Distribution of leverage in each of the six dominant strategies. Note the tendency for increased leverage in strategies of correlated debt–collateral pairs, especially those borrowing WETH (**ETH** → **ETH_{LST}**) and those looping yield-bearing USD-variants (**USD** → **USD**). We see the leverage being capped around 5× for strategies amplifying the directional exposure to tokens (**USD** → **Ex-USD**) due to the significantly lower liquidation thresholds.

Lastly, we examine the liquidations executed by the protocols throughout the observation period. While both interest rates and price fluctuations can negatively affect the funds posted by the borrower, liquidation remains a crucial risk factor. With liquidation penalties being currently around 6% for ETH positions, a liquidation could result in significant losses, considering that a full liquidation on a 10× leveraged position would imply a 54% loss of the borrower’s equity.




Fig. 6. The amount of debt that was repaid by liquidators to maintain overcollateralization. The liquidations are aggregated by month and are shown to be unevenly scattered across time. Liquidations are primarily associated with directional strategies, while tightly correlated tokens are largely absent.

Interestingly, we find that the strategies of tightly correlated pairs, even when they constitute the majority of the borrowing activity, have historically had negligible risk of liquidation. Instead, the vast majority of liquidations involved the strategies with directional exposure to volatile tokens, mostly those of WETH and WBTC.

5 Discussion

In this paper, we show how an increasingly large share of borrowed funds in decentralized lending markets consists of assets tightly correlated with their collateral. By examining address-level data, we deduce that agents increasingly use these markets to express their preference for high-leverage yield-maximization strategies, in which correlated yield-bearing and non-yield-bearing assets are looped to maximize the rewards obtained. This implies an increasing degree of sophistication among DeFi lenders and borrowers, many of whom now venture into trading strategies previously associated with hedge funds.

This structural shift has several notable implications. First, the most immediate consequence is the facilitation of substantially higher systemic leverage than previously feasible. By raising the liquidation threshold and enabling higher LTV for correlated token categories, lending markets implicitly facilitate a leveraging mechanism which, unlike that of perpetual futures (usually settled in USD-variant stablecoins), utilizes the underlying asset as collateral margin.

The elephant in the room is credit risk. Prior research warns that yield-bearing derivatives, such as staked or restaked ETH-variants, are not immune to abrupt dislocations. Episodes of rapid redemption demand in underlying yield-generating protocols such as Lido [16,14] or EigenLayer [2] coupled with the complex withdrawal queuing mechanism of the Ethereum staking set, can drain or reprice liquidity, causing sharp and sudden deviations from expected prices.

To mitigate systemic liquidations of correlated token pairs, Aave currently determines the price of derivative tokens through their underlying counterpart, rather than using direct price feeds of the derivative tokens themselves. The internal price of liquid staking tokens, such as wstETH, is therefore not determined by their market price. Instead, the price of wstETH is determined by the amount of ETH which can be redeemed at its issuer [1]. This solution effectively mitigates the risk of consuming pricing data from thinly traded markets, as this might invite strategic MEV attacks to force mass liquidations [18]. Nevertheless, the risk of mispricing collateral in the lending market relative to the price elsewhere imposes its own set of risks related to the solvency of the protocol.

As a result, lending markets explicitly rely on the assumption that users will eventually be able to withdraw their deposits from liquid staking protocols such as Lido. However, even when the staking protocols do remain solvent, the higher loan-to-value in conjunction with the new pricing mechanism can introduce new threats to the protocol if not correctly mitigated. One concern is that lending protocols can become exit liquidity for holders of liquid staking tokens in the event that redemption at the staking protocol ever were to stall [9,5]. If the

lending protocols operate on prices that are sufficiently above those in the secondary markets, one might borrow against one's staked tokens to access most of the liquidity, enabled by the higher LTV. This could drain the lending protocol of liquidity and raise the interest rates, forcing other borrowers to close their positions, either voluntarily or through liquidations, while the lack of liquidity could prevent lenders from withdrawing their tokens, effectively rendering lending protocols unable to operate until the redemption is resolved. Although this can be mitigated by pausing borrowing during such distressed situations, a mechanism referred to as "*Killswitch*", this requires correct execution, and one must consider its implications.

Another concern is how the new pricing mechanism can negatively affect the liquidation mechanism [9]. Since liquidations are executed by MEV arbitrageurs who are covering the unpaid debt by selling the collateral in the secondary markets, a dislocation between this price and the one used by the lending protocol could make the liquidation unprofitable unless accounted for.

Second, while previous literature has raised concerns around the emergence of competitive equilibria between on-chain lending markets and staking [6], the driving assumption here was that high demand for borrowing WETH for directional exposure might erode Ethereum's staking base, by drawing rational users away from staking when lending yields exceed staking rewards. However, our findings indicate that today, lending markets have had the opposite effect: throughout the entire period examined, nearly all borrowed WETH was ultimately used to stake directly or indirectly. The small fraction not staked corresponds to the liquidity buffer required for lender withdrawals. Even if this idle portion were staked, the total staked ETH supply would rise by less than two percent. This is far too little to pose a systemic risk.

Third, the data presented here also contribute toward explaining the changing interest rate environment within lending markets [7]. Historically, stablecoins commanded high lending rates because they were the preferred instruments for leveraging directional exposure to volatile assets such as WBTC and WETH [4]. Volatile tokens, by contrast, typically generated low returns for lenders due to abundant supply and limited borrowing demand. The introduction of yield-bearing tokens fundamentally altered this balance. Demand for borrowing has surged, as users pursue yield-arbitrage strategies rather than directional bets, pushing interest rates upward and contributing materially to the overall growth and attractiveness of lending protocols.

These findings suggest that correlated-asset strategies represent a meaningful shift in the economic foundations of DeFi lending, raising new questions about liquidation risk, protocol design, and borrower heterogeneity. This study is subject to several limitations. First, our dataset captures only address-level on-chain behavior within Aave v3 and SparkLend on Ethereum. Users may hedge, rebalance, or conduct offsetting trades on centralized exchanges, on other lending protocols, or on L2s, meaning that the leverage we observe may constitute an upper bound on the true economic exposure.

Second, our strategic inference necessarily depends on classification heuristics. Although the assets are grouped by correlation and functional similarity, users might pursue complex, multi-leg strategies that are not captured by these categories, introducing the possibility of misclassification. Third, as a result, the leverage metric abstracts away important behavioral heterogeneity such as net delta exposure, hedging, duration differences or portfolio-level risk management, limiting our ability to distinguish between structurally different trading styles that appear identical in aggregate. Fourth, the use of daily historical price sampling could miss short-lived price dislocations and intraday liquidation cascades, which may underestimate the true fragility of highly leveraged correlated positions.

Fifth, filtering out insolvent addresses and trivial balances introduces a survivorship bias toward more successful or sophisticated borrowers. Collectively, these limitations imply that while the overall structural trends we document in this paper are indicative of an increasing level of sophistication, the precise magnitudes of risk tolerance, leverage, and systemic vulnerabilities to the overall lending markets should be interpreted with caution.

6 Conclusion

Our findings demonstrate that lending activity is no longer driven by moderate leverage applied to speculative exposures or by simple liquidity extraction from volatile holdings. Instead, an increasingly large share of borrowing reflects highly leveraged positions constructed from pairs of strongly correlated assets, dominated by yield-bearing collateral against non-yield-bearing debt. This behavior seeks to capture interest rate spreads rather than directional exposure. As a result, correlated-asset strategies now constitute a dominant share of total market size. These developments carry meaningful implications for the stability, efficiency, and theoretical modeling of DeFi lending. The widespread use of correlated collateral-debt pairs enables higher leverage than traditional risk frameworks anticipate, potentially amplifying vulnerabilities during periods of liquidity stress or rapid repricing. At the same time, the rising demand for borrowing volatile assets can reshape interest rate dynamics, increasing yields for lenders and contributing to the continued growth of the lending protocols themselves. Future work should develop models that distinguish between correlated and uncorrelated leverage strategies and explore the systemic implications of yield-driven borrowing at scale, primarily with emphasis on modeling liquidations. Understanding these dynamics is essential for assessing both the resilience and the long-term trajectory of decentralized lending protocols.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

A Overview of Asset Categories

Mechanism category	Mechanisms	Tokens in dataset
ETH	ETH wrappers convert ETH into ERC-20 format for DeFi compatibility. Me- WETH chanically, these rely on mint/burn processes tied to locked collateral.	
Staked ETH	Tokens that represent staked ETH plus accrued staking rewards. Users de- ETHx, cbETH, osETH, rETH, wstETH posit ETH, the protocol stakes it via validators, and the LST then increases in value or supply to reflect earned yield. Some models rebase (ETHx, rETH) while others are value-accurring wrappers (wstETH). These give liquidity while maintaining staking exposure.	
Restaked ETH	Tokens representing ETH/LSTS that have been restaked in EigenLayer to ezzETH, rsETH, tETH, weETH provide crypto-economic security to additional services (AVSS). LRTs earn extra rewards on top of staking yield. Mechanically, they involve rehypoth- ecation of ETH trust assumptions across multiple services, adding reward layers but also layered slashing risk. We also include TETH in this category due to its higher yield from yield-strategies in addition to staking, despite it not being restaked.	
Stablecoins	Tokens designed to track USD value. Backing models differ: some are fully DAL, FRAX, GHO, LUSD, PYUSD, RLUSD, USDC, collateralized by fiat reserves (USDC, USDT) others use crypto-backed col- USDS, USDT, USDc, USDn, cryUSD, eUSDc, eUSDn lateral and stability modules (DAL, LUSD), and newer designs use on-chain yield-bearing assets or peg stabilizers (USDe). Their mechanism centers on mint/redeem flows, collateral buffers and peg arbitrage.	
Yield Stablecoins	Stablecoins that accrue yield. Additionally, we include principal tokens of PT - USDe - 25SEP2025, stablecoins which are created by splitting yield-bearing assets into PT (prin- PT - USDe - 27NOV2025, cipal) and YT (yield). PT represents the disbanded principal loanout, re- PT - eUSDe - 14AUG2025, deemable 1:1 at maturity. The mechanism behaves like a zero-coupon bond: PT - sUSDe - 25FEB2025, its price moves with interest rates and time to maturity. PT tokens do not PT - sUSDe - 27NOV2025, receive yield, only redemption value at expiry.	
BTC	BTC wrappers that lock Bitcoin on synthetic equivalents with custodians or smart contracts and mint a 1:1 representation on Ethereum.	
Others	Assets that do not follow any of the mechanisms above. This includes tokens AAVE, BAL, CRV, ENS, EURC, FXS, KNC, LDO, LINK, used to support parameters, control treasuries, direct emissions or capture MKR, ONE - INCH, RPL, SNX, STG, UNI, XAUT, GNO protocol fee flows in DAOs. They typically do not represent claims on un- derlying assets but provide influence over protocol evolution (AAVE, CRV, MKR) or utility rights (LINK for oracles). Mechanisms vary but revolve around on-chain voting and incentive alignment. Additionally, we include to- kenized gold (XAUT) and the stablecoin following the euro (EURC).	

B Assets Included in the Dataset and Variant Groupings

References

1. Aave: Oracle and pricing design documentation. Aave Developers. <https://docs.aave.com/developers/core-contracts/oracle> , accessed Jan. 2026.
2. Alexander, C.: Leveraged restaking of leveraged staking: What are the risks? SSRN 4840805 (2024).
3. Bartoletti, M., Chiang, J.-H.-y., Lluch-Lafuente, A.: SoK: Lending pools in decentralized finance. In: FC 2021 Workshops, LNCS 12675, 553–578 (2021).
4. Bertucci, C., Bertucci, L., Gontier-Delaunay, M., Guéant, O., Lesbre, M.: Agents’ behavior and interest rate model optimization in DeFi lending. Mathematical Finance (forthcoming 2025).
5. BGD (bgdlabs): Correlated-asset price oracle (CAPO). Aave Governance Forum proposal, Jan. 2024. <https://governance.aave.com/t/bgd-correlated-asset-price-oracle/16133> , accessed Jan. 2026.
6. Chitra, T.: Competitive equilibria between staking and on-chain lending. arXiv:2001.00919 (2020).
7. Cohen, S., N., Sánchez-Betancourt, L., Szpruch, L.: The economics of interest rate models in decentralised lending protocols. SSRN 4638390 (2023).
8. DeFiLlama: Lending protocols on Ethereum. Available at <https://defillama.com/protocols/lending/ethereum> , accessed Jan. 2026.
9. Gauntlet: Synchronicity Price Adapter “Killswitch” functionality for LST eMode. Aave Governance Forum proposal, May 2023. <https://governance.aave.com/t/gauntlet-synchronicity-price-adapter-killswitch-functionality-for-lst-emode/13224> , accessed Jan. 2026.
10. Heimbach, L., Huang, W.: DeFi leverage. BIS Working Paper No. 1171 (2024).
11. Luo, Y., Feng, Y., Xu, J., Tasca, P.: Piercing the veil of TVL: DeFi reappraised. arXiv:2404.11745 (2024).
12. Perez, D., Werner, S., M., Xu, J., Livshits, B.: Liquidations: DeFi on a knife-edge. In: FC 2021 (Part II), LNCS 12675, 457–476 (2021).
13. Qin, K., Zhou, L., Gamito, P., Jovanovic, P., Gervais, A.: An empirical study of DeFi liquidations: Incentives, risks, and instabilities. In: IMC 2021, ACM, 336–350 (2021).
14. Scharnowski, S., Jahanshahloo, H.: The economics of liquid staking derivatives: Basis determinants and price discovery. SSRN 4180341 (2023).
15. Soiman, F., Dumas, J.-G., Jimenez-Garcés, S.: What drives DeFi market returns? J. Int. Financial Markets Inst. Money **85**, 101786 (2023).
16. Xiong, X., Wang, Z., Chen, X., Knottenbelt, W., Huth, M.: Leverage staking with liquid staking derivatives (LSDs): Opportunities and risks. Preprint 2023-1842 (2024).
17. Xu, J., Feng, Y.: Reap the harvest on blockchain: A survey of yield farming protocols. IEEE Trans. Netw. Serv. Manag. **20**(1), 858–869 (2023).
18. Xu, J., Paruch, K., Cousaert, S., Feng, Y.: SoK: Decentralized exchanges (DEX) with automated market maker (AMM) protocols. ACM Comput. Surv. (to appear).