A Controlled Case Study of Design Trade-offs in DeFi Lending Protocols

Jan Findra[1], Zdenek Lapes[1], and Ivan Homoliak[1,2]

[1] Brno University of Technology, Faculty of Information Technology, Božetěchova 1/2, 612 00 Brno, Czech Republic
[2] Slovak University of Technology in Bratislava, Faculty of Informatics and Information Technologies, Ilkovičova 2, 842 16 Bratislava, Slovakia

Abstract. Decentralized lending protocols encode credit markets as smart contracts whose design choices affect gas cost, capital efficiency, liquidation behavior, and risk allocation. We compare Aave V3, SparkLend, Compound III, and Morpho Blue in matched Ethereum fork tests with wstETH collateral, USDC debt, and one fixed chain state. The tests measure call-path gas, annualized rates, thirty-day debt growth, maximum borrow capacity, and liquidation outcomes. It is important to note that the observed orderings for gas costs and rates are specific to this test artifact and could change under different collateral assets or utilization regimes. Within this setting, Morpho Blue has the lowest measured interaction call-path gas and highest borrow capacity, while shifting risk selection to isolated permissionless markets. Furthermore, Morpho Blue's evaluation uses a deployed market for rate measurements but a constructed market for liquidation testing. Aave V3 and SparkLend provide broader pooled-market abstractions with higher measured call-path gas. Compound III simplifies borrowing around one base asset per market; its low-gas `absorb` step is offset by a separate `buyCollateral` liquidation path. We provide a Foundry-based Ethereum-fork test suite for reproducing these measurements and comparing lending-design trade-offs.

Keywords: decentralized finance · lending protocols · liquidation · smart contracts · call-path gas measurement · capital efficiency

1 Introduction

Decentralized finance (DeFi) replaces parts of financial intermediation with deterministic programs executed on public blockchains [1] [2]. Lending is a core DeFi primitive because it supports leverage, collateralized liquidity, interest-bearing deposits, stablecoin demand, and liquidation markets. Unlike traditional credit markets, most on-chain lending markets cannot rely on legal enforcement against defaulting borrowers. They therefore use over-collateralization, price oracles, liquidation incentives, and programmatic interest-rate models to maintain solvency [3] [4] [5]. These protocols encode credit-market rules directly in executable state: users supply collateral, open debt, accrue interest as market

state changes, and become liquidatable when a collateral constraint is violated. Small differences in these rules affect user cost, liquidity efficiency, liquidation incentives, governance risk, and the attack surface.

This paper compares four deployed Ethereum lending protocols under one executable market scenario. Aave V3 represents shared pooled liquidity [6] [7]. SparkLend is Aave-derived but governed and configured for the Sky ecosystem [8]. Compound III, also called Comet, uses a single borrowable base asset per market [9] [10] [11]. Morpho Blue uses immutable isolated markets with permissionless risk parameter selection and a singleton contract architecture [12] [13].

The research questions are:

1. **RQ1:** How do matched wstETH/USDC lending operations differ in measured call-path gas across the four protocol designs?
2. **RQ2:** How do observed rates, thirty-day debt growth, and maximum borrow capacity differ under a common wstETH/USDC scenario?
3. **RQ3:** What security and usability trade-offs are exposed by this controlled comparison of pooled liquidity, single-base-asset markets, and isolated permissionless markets?

We provide a reproducible empirical case study that compares deployed lending designs under matched conditions. The experiment fixes the collateral asset, debt asset, borrower balance, time advance, and liquidation shock where possible. It records quantitative metrics and integration observations.

2 Related Work

Gudgeon et al. introduced protocols for loanable funds and analyzed how DeFi lending protocols set interest rates in response to liquidity and utilization [3]. Xu and Vadgama discuss the shift from traditional lending to over-collateralized DeFi lending [4]. Bartoletti et al. systematize lending pools as parametric incentive mechanisms, formalize lending-pool behavior, and model lending protocols as economic mechanisms with strategic and adversarial dynamics [5] [14] [15]. Liquidation work shows that liquidable positions can be resolved quickly, that liquidator incentives transfer discounted collateral away from borrowers, and that close-factor rules can amplify stress under adverse collateral conditions [16] [17] [18]. Our comparison complements this work by executing matched scenarios across four current protocol designs.

Prior security studies attribute many DeFi failures to interactions among economic design, permissionless execution, oracle dependence, transaction ordering, and composability. Werner et al. distinguish technical security from economic security in DeFi [2]. Daian et al. show that transaction ordering and MEV create competitive fee dynamics in smart-contract markets [19]. Zhou et al. systematize DeFi attacks, including price-oracle attacks and permissionless interactions [20]. Oracle studies explain why off-chain data feeds create trust and manipulation problems for smart contracts, including lending markets [21] [22].

Aave V3, Compound III, Morpho Blue, and SparkLend documentation define the mechanisms evaluated here: shared pools and Aave-style risk parameters, a single borrowable base asset, immutable isolated markets with permissionless risk management, and Spark-specific Aave-derived configuration [6] [7] [9] [10] [11] [12] [23] [8].

3 Protocol Design Background

The compared protocols differ in liquidity organization, risk-parameter control, and liquidation design. Aave V3 and SparkLend use Aave-style shared pools and health-factor liquidation, with Spark-specific governance and configuration [7] [24] [8]. Compound III uses one borrowable base asset per market while treating other assets as collateral only [9] [10]. Morpho Blue defines each isolated market by loan token, collateral token, oracle, interest-rate model, and liquidation loan-to-value parameter [23] [12].

4 Methodology

The measurement artifact is implemented as Foundry tests executed against an Ethereum mainnet fork [25] [26]. The tests import only the interface files needed to call the deployed contracts. A shared base test contract defines the common assets and constants. The collateral asset is wrapped staked Ether, wstETH, and the debt asset is USDC. The tested account receives one wstETH of collateral. The Aave V3 oracle is used as the reference price source for common calculations [7]. The normal-borrow scenario borrows 60% of the collateral value. The liquidation scenario opens a larger position at 75% of the collateral value and then applies

Table 1. High-level design comparison of the evaluated lending protocols.

Attribute	Aave V3	SparkLend	Compound III	Morpho Blue
Liquidity	Shared multi-asset pool	Aave-derived shared pool	Single base-asset market	Isolated loan/ collateral pair
Risk control	Governance-set reserves	Sky-governed configuration	Market collateral factors	Creator selects LLTV, oracle, and IRM
Oracle model	Protocol oracle	Protocol plus Spark-specific components	Protocol price feeds	Oracle chosen per market
Upgrade model	Upgradeable protocol system	Upgradeable Aave-derived system	Upgradeable Comet market	Immutable core protocol
Integration benefit	Shared liquidity and established tooling	Aave-like interface in Sky ecosystem	Simpler base-asset borrowing	Low call-path gas and flexible markets

a 20% wstETH price reduction through mocked oracle behavior or an equivalent protocol-specific price-feed adjustment.

For Morpho Blue, the deployed Ethereum wstETH/USDC market is used for rate, utilization, debt-accrual, capital-efficiency, and normal interaction measurements. The deployed Morpho market identifier is `0xb323495f7e4148be5643a4ea4a8221eef163e4bccfdedc2a6f4696baacbc86cc`. The test loads the market parameters from the deployed Morpho singleton through this identifier, preserving the market's existing oracle, interest-rate model, LLTV, liquidity, and utilization state. The forced liquidation scenario uses a separate constructed market with the same token pair, adaptive interest-rate model, and 86% LLTV, but with a mock oracle. This split keeps Morpho's rate and accrual measurements tied to a deployed market while retaining a deterministic 20% price shock for liquidation mechanics.

The experiment contains four scenarios.

1. **Full flow:** supply collateral, borrow USDC, advance time by 30 days, repay the debt, and withdraw collateral. Call-path gas is logged for supply, borrow, repay, and withdraw.
2. **Interest accrual:** open a borrow position and advance the fork one day at a time for 30 days. The debt balance is recorded each day.
3. **Liquidation simulation:** open a near-threshold borrow, reduce the collateral price by 20%, verify that the account becomes liquidatable, execute liquidation, and record liquidation call-path gas, repaid or absorbed USDC debt, and collateral received or absorbed. For Compound III, the artifact records both `absorb` and `buyCollateral` because liquidation settlement is split across those calls.
4. **Capital efficiency:** read or derive loan-to-value limits, liquidation thresholds, liquidation bonus where available, and the maximum USDC borrowable for one wstETH.

Fork testing keeps deployed contract behavior and selected mainnet state. The tests fix account balances, time advances, and price shocks so each protocol is exercised under the same local execution environment. Gas is reported in gas units because Ethereum defines gas as the unit of computational effort required to execute operations, and the fiat cost of a transaction also depends on the gas price at inclusion time [27]. The gas counters bracket individual protocol calls inside Foundry tests by reading `gasleft()` before and after the call. These values

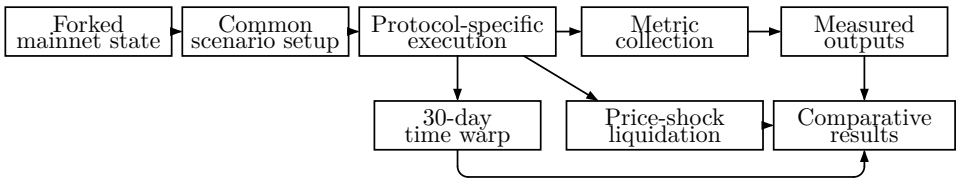


Fig. 1. Experimental workflow. Each scenario produces comparable measurements for call-path gas, debt accrual, liquidation behavior, and capital efficiency.

are measured call-path gas, not complete user transaction gas. They exclude the outer transaction envelope, approval-call gas, wallet flow, transaction calldata pricing, live mempool competition, MEV, liquidator bidding, L1 congestion, L2 fee markets, and real oracle update timing.

The reported transaction counts are workflow assumptions rather than Foundry transaction counts. With no pre-existing token allowances, the modeled full flow requires six user-submitted transactions for each protocol: two approvals plus supply, borrow, repay, and withdraw-style calls. Existing allowances would reduce this count. The gas tables report only the protocol operation calls listed in the table, not approval calls or complete user-cost estimates.

4.1 Mathematical Definitions of Measured Quantities

Let C be the supplied wstETH amount, P_c the collateral price in USDC, and $V_c = CP_c$ the collateral value. The normal borrow and liquidation-scenario borrow amounts are fixed as $B_{\text{normal}} = 0.60V_c$ and $B_{\text{liq}} = 0.75V_c$.

For each protocol p, thirty-day debt growth is measured from protocol state as $\Delta D_{p(T)} = D_{p(T)} - D_{p(0)}$, with $T = 30$ days. Maximum borrow capacity is reported as $B_{\max} = V_c\lambda_b$, where λ_b is the protocol-specific borrow factor. For Aave V3 and SparkLend this is the reserve LTV, for Compound III it is the borrow collateral factor, and for Morpho Blue it is the market LLTV [7] [9] [23].

The liquidation experiment applies a 20% collateral-price shock, so $V_c' = 0.80V_c$. For Aave-style protocols, liquidation is possible when $\text{HF} = V_c\frac{\lambda_L}{D} < 1$ [24]. For Morpho Blue, liquidation is possible when $\frac{D}{V_c} > \text{LLTV}$ [23]. For Compound III, the analogous condition is that debt exceeds the value allowed by the liquidation collateral factor [9].

For Aave-style liquidation bonuses, a liquidator repaying R debt receives collateral with oracle value $V_{\text{seized}} = R\beta_L$. The pre-cost spread is therefore $\Pi_g = R(\beta_L - 1)$. If available post-shock collateral value is limited to V_c', the maximum repay supported by the remaining collateral is $R_{\max} = \frac{V_c'}{\beta_L}$. This explains the collateral-limited SparkLend result: using the SparkLend values in Table 5, the post-shock collateral value divided by the 1.07 liquidation-bonus factor is below the liquidation-scenario debt. The liquidation consumes essentially all collateral and leaves residual debt.

For Morpho Blue, the liquidation incentive is determined by the Liquidation Incentive Factor [23]:

$$\text{LIF} = \min\left(M, \frac{1}{\beta\ \text{LLTV} + (1-\beta)}\right), \quad \beta = 0.3, \quad M = 1.15. \tag{1}$$

Compound III reports the `buyCollateral` discount spread separately because `absorb` and discounted collateral purchase are separate calls [9] [11]. Measured call-path gas is $G_{\text{call}} = g_{\text{before}} - g_{\text{after}}$. The reported gas values exclude approval calls, transaction-envelope costs, gas-price volatility, routing, slippage, and MEV competition [27] [19].

Table 2. Measured call-path gas for matched protocol operations in the experimental artifact. Approval calls and transaction-envelope costs are excluded.

Operation	Aave V3	Compound III	Morpho Blue	SparkLend
Supply	189,455	156,850	79,729	208,522
Borrow	212,086	184,117	173,692	236,570
Repay	52,810	49,277	19,960	61,097
Withdraw	66,470	80,375	35,633	71,754
Full-flow total	520,821	470,619	309,014	577,943
Liquidation / absorb	174,374	90,160	73,778	191,231
buyCollateral	N/A	98,178	N/A	N/A
Liquidation total	174,374	188,338	73,778	191,231
Full flow + liquidation	695,195	658,957	382,792	769,174

5 Results

5.1 Call-Path Gas

Table 2 gives the measured call-path gas for matched protocol operations. For Morpho Blue, supply, borrow, repay, withdraw, rate, and accrual use the deployed wstETH/USDC market; liquidation uses the constructed mock-oracle market described in the methodology. In these paths, Morpho Blue uses the least gas. Its measured full-flow call total is 309,014, compared with 470,619 for Compound III, 520,821 for Aave V3, and 577,943 for SparkLend. That ordering is unchanged when the constructed Morpho liquidation row is added. For Compound III, the table separates the low-gas `absorb` call from the later discounted `buyCollateral` purchase and reports their sum as the full liquidation call path.

Morpho Blue's low measured gas is consistent with its minimal singleton architecture, although the experiment does not attribute the difference to individual implementation-level causes [12] [13]. Compound III also benefits from the single-base-asset design, especially in the `absorb` phase, but the full liquidation economic path includes the separate `buyCollateral` call. The difference between Aave V3 and SparkLend is smaller conceptually but visible empirically. Since SparkLend inherits the Aave V3 design but modifies ecosystem-specific components, the higher measured call-path gas in this test is an artifact-level observation rather than a universal claim across all Spark deployments or assets.

5.2 Rates and Interest Accrual

Table 3 reports the APY values extracted by the tests. SparkLend has the highest borrow APY in this scenario, followed by Aave V3, Compound III, and Morpho Blue. For Morpho Blue, the borrow rate is read from the deployed market's IRM after accruing interest. Its supply APY is derived from borrow rate, utilization, and market fee rather than read from a single pool-wide supply-rate field. Aave V3 and SparkLend expose liquidity rates directly, and Compound III exposes

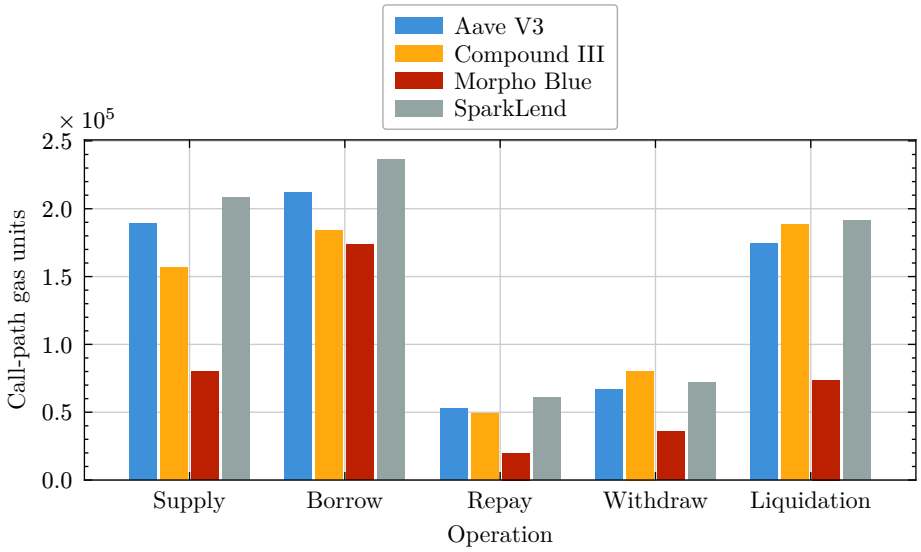


Fig. 2. Measured call-path gas by operation for the matched protocol paths. Morpho Blue has the lowest measured call path in every operation in this artifact.

a Comet supply rate. The Morpho Blue rates are measured from the deployed wstETH/USDC market rather than from the constructed liquidation market.

The observed 30-day debt growth is highest for SparkLend. Aave V3, Morpho Blue, and Compound III are closer together, and their middle ordering differs from the point APY row because the debt-growth test observes protocol-specific accrual along the simulated borrow path. Borrowing costs remain state-dependent because utilization, governance parameters, and market configuration change over time.

Table 3. Observed rates and thirty-day debt growth for the matched borrowing scenario.

Metric	Aave V3	Compound III	Morpho Blue	SparkLend
Borrow APY	4.058%	3.781%	3.746%	4.664%
Supply APY	3.330%	2.907%	3.370% (derived)	4.126%
Debt at day 0	1,742.285 USDC	1,742.285 USDC	1,742.285 USDC	1,742.285 USDC
Debt at day 30	1,747.990 USDC	1,747.599 USDC	1,747.744 USDC	1,748.826 USDC
Thirty-day debt growth	5.706 USDC	5.314 USDC	5.460 USDC	6.541 USDC

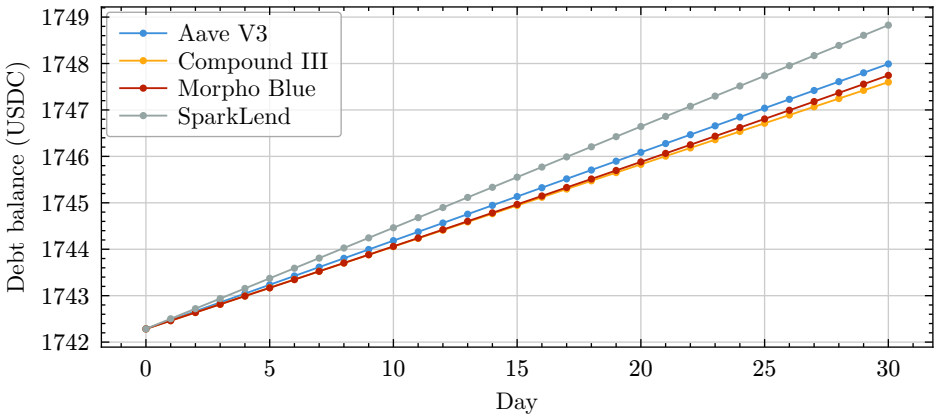


Fig. 3. Thirty-day debt accrual under the matched borrowing scenario.

5.3 Capital Efficiency

Capital efficiency is measured by the borrow loan-to-value or borrow collateral factor where available, liquidation threshold or LLTV, liquidation incentive mechanism, and maximum USDC borrowable against one wstETH. Morpho Blue has the highest borrow capacity in the tested configuration, followed by SparkLend, Compound III, and Aave V3. Aave V3 and SparkLend expose Aave-style liquidation bonuses in the measured data. Compound III and Morpho Blue use different incentive mechanisms, so Table 4 reports the mechanism rather than encoding them as zero bonuses.

Table 4. Capital-efficiency metrics for one wstETH collateral in the tested configuration.

Metric	Aave V3	Compound III	Morpho Blue	SparkLend
Borrow factor	78.5%	82.0%	N/A	83.0%
Liq. threshold / LLTV	81.0%	86.0%	86.0%	84.0%
Liquidation incentive	6.0% bonus	`buyCollateral` discount	LIF(LLTV)	7.0% bonus
Max borrow	2,279.49	2,375.92	2,486.04	2,405.08

Note. Compound III uses a two-stage liquidation path: `absorb` transfers the position to the protocol and `buyCollateral` sells collateral at a discount. Morpho Blue uses a liquidation incentive factor derived from LLTV [9] [11] [23].

Morpho Blue does not have a separate borrow LTV below the liquidation threshold in the same sense as Aave-derived or Compound III markets. The reported 86.0% value is the market LLTV, which is both the maximum safe LTV boundary and the liquidation trigger. Although the Borrow LTV row is marked N/A for Morpho

Blue, the maximum-borrow calculation uses the market LLTV as the boundary. This differs from Aave-derived and Compound III markets, where borrow capacity and liquidation eligibility are represented by separate parameters.

Higher borrow capacity lets leverage-seeking borrowers open larger positions and leaves less room before liquidation. Capital efficiency therefore depends on oracle quality, collateral volatility, liquidation speed, and market liquidity. The Morpho Blue measurements show this trade-off directly. The deployed Morpho market provides the capital-efficiency, rate, and normal-flow measurements. The constructed Morpho market isolates the liquidation mechanism under a deterministic mock-oracle shock. In both cases, users must trust the selected market parameters and oracle.

5.4 Liquidation Behavior

The liquidation scenario opens a 75% collateral-value borrow and applies a 20% collateral price drop. All protocols become liquidatable under the test assumptions. The first liquidation table reports collateral received by the caller, debt repaid or absorbed, and post-shock caller spread. The second reports call-path gas for the liquidation path. The Morpho Blue liquidation row comes from the constructed mock-oracle market described in the methodology. For Aave V3, SparkLend, and Morpho Blue, spread is the post-shock oracle value of collateral received by the liquidator minus debt repaid. For Compound III, `absorb` transfers the underwater account's collateral and debt to the protocol, while the caller-facing economic spread comes from the separate discounted `buyCollateral` purchase valued at the post-shock price [9] [11].

Table 5. Liquidation economic outcomes after a 20% collateral price reduction. Spread values received collateral at the post-shock oracle price and exclude gas price, slippage, and MEV effects; the Compound III spread is the buyCollateral discount spread.

Metric	Aave V3	Compound III	Morpho Blue	SparkLend
Initial wstETH	1.000	1.000	1.000	1.000
Borrowed (USDC)	2,177.86	2,177.86	2,177.86	2,177.86
Caller wstETH	0.9880	0.9908	0.9786	0.9935
Debt repaid / absorbed	2,177.86	2,177.86	2,177.86	2,171.07
Absorbed value	N/A	2,323.04	N/A	N/A
Post-shock spread	117.31	123.85	95.47	136.77

Table 6. Liquidation call-path gas after a 20% collateral price reduction. Compound III separates `absorb` and `buyCollateral`; total gas reports their sum.

Operation	Aave V3	Compound III	Morpho Blue	SparkLend
Liquidation / absorb	174,374	90,160	73,778	191,231
buyCollateral	N/A	98,178	N/A	N/A
Liquidation call-path gas	174,374	188,338	73,778	191,231

SparkLend repays less than the initial liquidation-scenario debt because the 20% price shock combined with the 7.0% liquidation bonus makes the position collateral-limited. With one wstETH of collateral, the protocol cannot repay the full debt while also awarding the liquidation bonus, so the liquidation consumes essentially all available collateral and leaves the repaid amount below the original borrow.

Caller spread is a pre-cost protocol-level measure computed after the price shock. It excludes gas price, priority fees, routing, slippage, inventory cost, and MEV competition. These exclusions matter because priority gas auctions, MEV, and liquidation spirals can change realized liquidator returns and the systemic effect of liquidation opportunities [19] [18]. Within the controlled price-shock scenario, the protocols differ in collateral received by the caller, debt repaid or absorbed, and liquidation call-path gas.

6 Discussion

In the wstETH/USDC setting, the measurements show trade-offs that align with the protocols' architectural choices. The deployed Morpho market provides the rate, capital-efficiency, and normal-flow measurements, while the constructed

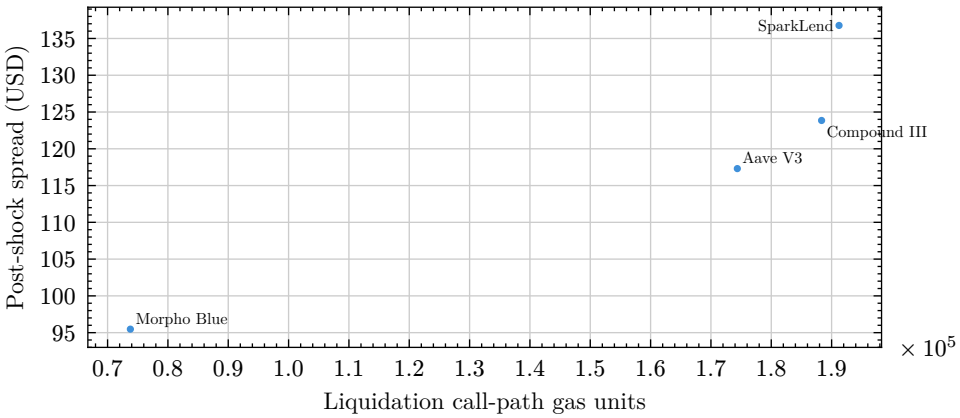


Fig. 4. Liquidation call-path gas versus post-shock caller spread in the controlled price-shock scenario. The Compound III point uses buyCollateral discount spread.

Morpho market isolates liquidation under a deterministic price shock. Morpho Blue has the lowest measured call-path gas in the tested artifact and allows the highest borrow amount against one wstETH. Its architecture separates the lending primitive from the risk-management layer more explicitly than the other evaluated protocols. This design gives expert users and protocol builders more control, but passive users must select among markets with different oracles, liquidation loan-to-value values, and interest-rate models.

Aave V3 and SparkLend expose a broader pooled-market abstraction in which multiple assets share an account-level risk model. This design supports multi-asset account management, but it also introduces additional governance, oracle, and parameter-management surfaces. Shared pools also create risk-management challenges: a parameter change, liquidity shock, or oracle issue in one listed asset can affect pool safety. System-level DeFi analyses similarly emphasize interconnectedness, leverage, and limited shock absorption as risk channels [28] [29]. Governance and upgrades are therefore not implementation details; they are part of the economic design.

Compound III occupies an intermediate position in the evaluated design space. By allowing one borrowable base asset and treating other assets only as collateral, it reduces some risk paths common to multi-asset pools. It also shows relatively low measured `absorb` call-path gas in the test, while the full liquidation path includes the later discounted collateral-purchase call. The trade-off is less asset-level flexibility. A user who wants to borrow different assets needs different Comet markets, and each market has its own base-asset dynamics.

The measurements show that execution cost, collateral constraints, and liquidation incentives interact even in one market state. Lower liquidation call-path gas reduces one component of liquidation execution cost and can make smaller liquidation opportunities economically viable when other costs are fixed. A protocol with higher LTV can improve capital efficiency but reduce the liquidation buffer. A protocol with permissionless markets can diversify risk profiles but also increase user due-diligence requirements. Global ranking would require additional assets, block heights, chains, utilization regimes, and liquidation conditions.

The artifact also exposed integration and security differences relevant to RQ3. Aave V3 and SparkLend were the most similar from the developer perspective because both use Aave-style pool interactions. Compound III required adapting to the Comet model, especially the distinction between the base asset and collateral-only assets [9] [10]. Morpho Blue required a market-parameter object for every interaction; the deployed-market test loads that object from the market identifier, while the constructed liquidation test creates a local fork market with a mock oracle. Oracle dependence remains a common risk channel. Aave V3, SparkLend, and Compound III use protocol-level price-feed assumptions, while Morpho Blue allows oracle choice per market [30] [23]. Prior oracle-attack work places oracle design inside the protocol security model rather than outside it [21] [22]. Upgradeability adds a related trade-off: governance-controlled systems can patch and reconfigure, while immutable systems reduce intervention risk but cannot patch

the core protocol if a flaw is found [12] [6]. All tested protocols require token approvals for at least some actions, so approvals remain part of the real lending workflow even though they are excluded from the call-path gas tables.

7 Threats to Validity

The controlled design improves internal comparability but limits external validity. The experiment uses one collateral asset, one borrowed asset, and one chain state. Different token implementations, collateral volatilities, oracle configurations, reserve parameters, and utilization levels can change rates, call-path gas, risk parameters, and liquidation outcomes. The measured ordering may therefore change across assets, block heights, chains, and utilization regimes.

The execution model omits mainnet transaction competition, MEV, liquidity routing, approval transaction costs, transaction-envelope costs, calldata pricing, and gas-price volatility. Prior MEV work shows that transaction ordering and priority-fee competition can materially affect on-chain opportunity capture, and Ethereum documentation separates gas usage from the market price paid per gas unit [19] [27]. Mocked price drops isolate liquidation mechanics but do not reproduce oracle update delay, exchange-price manipulation, or cross-market contagion.

Construct validity is limited by metric differences across protocols. Aave-derived reserve rates, a Compound III base-market rate, and a Morpho isolated-market rate reflect different market structures. Morpho Blue uses a deployed market for rate and normal-flow measurements but a constructed market for the forced liquidation row, so those Morpho measurements should not be collapsed into a single deployed-market claim. The implementation uses selected interfaces copied from official repositories rather than full protocol source trees. Extensions should check those interfaces against current protocol repositories before adding new measurements.

An expanded benchmark should evaluate more collateral and debt pairs, test multiple block heights and utilization regimes, include L2 deployments, replay historical liquidation events, and compute liquidator profit after estimated gas prices and swap slippage.

Furthermore, some of the reported differences may reflect protocol configuration choices rather than purely architectural design decisions. Market-specific parameters such as liquidation thresholds, oracle selection, and interest-rate models inevitably influence the measured outcomes alongside the underlying protocol architecture.

8 Conclusion

This paper presented a measurement case study of Aave V3, SparkLend, Compound III, and Morpho Blue under one controlled wstETH/USDC setting. The results identify design positions across flexibility, capital efficiency, execution cost,

governance dependence, and user risk selection. Within this specific artifact, Morpho Blue minimizes measured call-path gas and maximizes borrow capacity in the tested configuration, while shifting more risk-selection responsibility to market creators and users. Compound III simplifies borrowing through a single base asset, and Aave V3 and SparkLend provide broader pooled-market abstractions with higher measured call-path gas in this artifact. The results caution against APY-only comparison because borrow cost, execution cost, liquidation path, oracle dependence, market depth, and governance model all affect protocol behavior. Furthermore, the gas and rate orderings observed here are specific to the tested conditions and could change under different scenarios.

Artifact Availability

The Foundry-based Ethereum-fork test suite and measurement artifact used to generate the results in this paper are open-source and available at https://github.com/zdeneklapes/bsd/tree/main/code/xfindr01.

Acknowledgement

This work was supported by the Science Grant Agency VEGA under project 1/0771/26, and by the Brno University of Technology internal project FIT-S-23-8151.

Declaration of generative AI use

The authors used OpenAI ChatGPT to support drafting, language editing, structural revision, and citation checking. All AI-assisted output was reviewed and verified by the authors, who take full responsibility for the final content. No AI tool is listed as an author, and no AI-generated experimental data or images are included.

Bibliography

1. Schär, F.: Decentralized Finance: On Blockchain- and Smart Contract-Based Financial Markets. Federal Reserve Bank of St. Louis Review. 103, 153–174 (2021). https://doi.org/10.20955/r.103.153-74.
2. Werner, S.M., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., Knottenbelt, W.J.: SoK: Decentralized Finance (DeFi). In: Proceedings of the 4th ACM Conference on Advances in Financial Technologies. pp. 30–46 (2022). https://doi.org/10.1145/3558535.3559780.
3. Gudgeon, L., Werner, S., Perez, D., Knottenbelt, W.J.: DeFi Protocols for Loanable Funds: Interest Rates, Liquidity and Market Efficiency. arXiv preprint arXiv:2006.13922. (2020).
4. Xu, J., Vadgama, N.: From Banks to DeFi: The Evolution of the Lending Market. arXiv preprint arXiv:2104.00970. (2021).

5. Bartoletti, M., Chiang, J.H.-y., Lluch-Lafuente, A.: SoK: Lending Pools in Decentralized Finance. In: Financial Cryptography and Data Security. pp. 553–578 (2021). https://doi.org/10.1007/978-3-662-63958-0_40.
6. Frangella, E., Pons, E.B., Miroshkin, A., Contributors, A.: Aave V3 Technical Paper, https://github.com/aave/aave-v3-core/blob/master/techpaper/Aave_V3_Technical_Paper.pdf.
7. Aave: Aave V3 Documentation, https://aave.com/docs/aave-v3/overview, last accessed 2026/06/20.
8. Spark Protocol: SparkLend Security Audits and Aave V3 Codebase Notes, https://docs.spark.fi/dev/security/security-and-audits, last accessed 2026/06/20.
9. Compound Finance: Compound III Documentation, https://docs.compound.finance/, last accessed 2026/06/20.
10. Compound Finance: Collateral and Borrowing, https://docs.compound.finance/collateral-and-borrowing/, last accessed 2026/06/20.
11. OpenZeppelin: Compound III Audit, https://www.openzeppelin.com/news/compound-iii-audit.
12. Delaunay, M.G., Frambot, P., Garchery, Q., Lesbre, M.: Morpho Blue Whitepaper, https://github.com/morpho-org/morpho-blue/blob/main/morpho-blue-whitepaper.pdf.
13. Morpho Labs: morpho-blue: Morpho Blue Protocol Repository, https://github.com/morpho-org/morpho-blue, last accessed 2026/06/20.
14. Bartoletti, M., Chiang, J.H.-y., Junttila, T., Lluch-Lafuente, A., Mirelli, M., Vandin, A.: Formal Analysis of Lending Pools in Decentralized Finance. arXiv preprint arXiv:2206.01333. (2022).
15. Bartoletti, M., Lipparini, E.: A Theory of Lending Protocols in DeFi. arXiv preprint arXiv:2506.15295. (2025).
16. Perez, D., Werner, S.M., Xu, J., Livshits, B.: Liquidations: DeFi on a Knife-edge. In: Financial Cryptography and Data Security. pp. 457–476 (2021). https://doi.org/10.1007/978-3-662-64331-0_24.
17. Qin, K., Zhou, L., Gamito, P., Jovanovic, P., Gervais, A.: An Empirical Study of DeFi Liquidations: Incentives, Risks, and Instabilities. arXiv preprint arXiv:2106.06389. (2021).
18. Warmuz, J., Chaudhary, A., Pinna, D.: Toxic Liquidation Spirals. arXiv preprint arXiv:2212.07306. (2022).
19. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A.: Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges. In: 2020 IEEE Symposium on Security and Privacy (SP). pp. 910–927 (2020). https://doi.org/10.1109/SP40000.2020.00040.
20. Zhou, L., Xiong, X., Ernstberger, J., Chaliasos, S., Wang, Z., Wang, Y., Qin, K., Wattenhofer, R., Song, D., Gervais, A.: SoK: Decentralized Finance (DeFi) Attacks. In: 2023 IEEE Symposium on Security and Privacy (SP). pp. 2444–2461 (2023). https://doi.org/10.1109/SP46215.2023.10179435.
21. Eskandari, S., Salehi, M., Gu, W.C., Clark, J.: SoK: Oracles from the Ground Truth to Market Manipulation. arXiv preprint arXiv:2106.00667. (2021).
22. Aspembitova, A.T., Bentley, M.A.: Oracles in Decentralized Finance: Attack Costs, Profits and Mitigation Measures. Entropy. 25, 60 (2023). https://doi.org/10.3390/e25010060.
23. Morpho Association: Variable Rate Market (Morpho Blue), https://docs.morpho.org/learn/concepts/market/, last accessed 2026/06/20.
24. Aave: Health Factor and Liquidations, https://aave.com/help/borrowing/liquidations, last accessed 2026/06/20.

25. Foundry Contributors: Foundry Book, https://book.getfoundry.sh/, last accessed 2026/06/20.
26. Foundry Contributors: Fork Testing, https://www.getfoundry.sh/guides/fork-testing, last accessed 2026/06/20.
27. Ethereum Foundation: Gas and Fees: Technical Overview, https://ethereum.org/developers/docs/gas/, last accessed 2026/06/20.
28. Aramonte, S., Huang, W., Schrimpf, A.: DeFi Risks and the Decentralisation Illusion. BIS Quarterly Review. 21–36 (2021).
29. Financial Stability Board: The Financial Stability Risks of Decentralised Finance, https://www.fsb.org/2023/02/the-financial-stability-risks-of-decentralised-finance/.
30. Chainlink Labs: Price Feeds, https://docs.chain.link/data-feeds/price-feeds, last accessed 2026/06/20.