OPEN INTEREST RESEARCH · PAPER NO. 1 · NOVEMBER 2026

Parametric Cover on DeFi Borrowing Costs: An Out-of-Sample Study of USDC Variable Rates on Compound and Aave, 2019–2026

Open Interest research desk · Working paper, not peer reviewed · Data through Nov 15, 2026

ABSTRACT. We rebuild a daily USDC variable borrow rate for Ethereum mainnet from June 2019 to November 2026. The rate comes straight from the interest indexes of Compound v2, Aave v2 and Aave v3, read at 2,848 archive blocks. On it we price, day by day and strictly out of sample, the 30-day rate-cap layer Open Interest sells: a 10-point layer above a cap set near the current 30-day average. Pricing uses the production engine, which takes a trailing-year burn-cost distribution, adds a stress month, applies a Wang distortion and charges for capital. Suppose 5% of all outstanding USDC borrowing had been insured, about $112m in force. At the default cap the pool would have written $2.72m of premium a year and paid $0.24m in claims: a loss ratio of 8.9%, with a 90% block-bootstrap interval of 5%–15%. Capital is never at risk of insolvency because every policy is fully collateralized, and the pool lost money in no calendar year. Seen from the buyer, the cover is expensive. It costs about 2.43% of notional a year to protect a rate that averaged 6.3%, and it barely reduces the variability of 30-day borrowing cost. The price breaks down into a 2.4× risk loading on the burn cost and a burn cost that overstates realized claims by 2.9×. We trace the second factor to a "spike hangover" in trailing-window rating. We then report a pricing frontier for recalibration.

KEYWORDS parametric insurance · interest-rate caps · decentralized lending · Wang transform · burn-cost rating · backtesting

1. Introduction

Borrowers on pooled lending protocols such as Compound and Aave pay a variable rate set every block by a utilization curve (Leshner and Hayes, 2019; Aave, 2020). Rates are low and stable most of the time, then jump when demand to borrow spikes: in leveraged bull markets, during liquidity-mining programs, or when a stablecoin briefly loses its peg (Gudgeon et al., 2020). Treasury desks, basis traders and leveraged stakers all carry this risk. Traditional finance hedges it with caps or swaps (Black, 1976). On-chain there is almost no market for that hedge.

Open Interest sells a simple parametric cap. For a fixed 30-day period, the buyer is paid if the period's average borrow rate ends above a chosen level. The payout grows linearly up to a fixed width and then stops. Settlement uses a public index, so no claim is filed. This paper asks four questions about that product, using the longest on-chain record available:

(i) What does the USDC borrow-rate distribution look like over the full history of variable-rate lending? (ii) How would the production pricing engine have performed out of sample? (iii) At a realistic take-up, what premiums, claims and capital would have resulted? (iv) Why does the engine land where it does, and how might it be recalibrated?

On the "since 2017" framing: on-chain variable-rate lending of USDC started with Compound v2 in May 2019. ETHLend (2017) matched fixed-term peer-to-peer loans and Compound v1 (September 2018) had no USDC market, so neither produces a comparable index. Our sample therefore starts in June 2019, the first month in which all the data we need can be read.

2. The contract and its price

2.1 Payout

Let $\bar{R}$ be the arithmetic mean of the daily borrow APY over the 30 days of cover, in percent. The cap K is the attachment point and w the layer width, fixed at 10 percentage points. Per dollar of insured notional the payout is

$$X = \min(\max(\bar{R} - K, 0), w) / 100 \times 30/365 \tag{1}$$

so the limit is $L = w/100 \times 30/365 = 0.822\%$ of notional. In reinsurance language the cover is a "10 points xs K" layer. In options language it is a call spread on the average rate, struck at K and $K + w$. The pool contract reserves $L - P$ of underwriter capital per policy, so the most a policy can lose is what is reserved for it.

2.2 Caps

Four caps are offered. With m the current 30-day average and σ the standard deviation of rolling 30-day averages over the trailing year, they are $K_{0..3} = m + \{0, ½, 1, 1½\}\cdot\sigma$, each rounded up to a 0.25-point grid. K_1 is the default the interface highlights, and the rest of the paper calls it "the default cap". Over our sample K_1 averaged 6.98%, sitting 1.22% above the trailing 30-day average.

2.3 Premium

The engine forms an empirical loss distribution from the payouts every rolling 30-day window of the trailing 365 days would have produced. It gives that sample weight 23/24 and puts weight 1/24 on a single stress month $s = \min(\text{worst day}, \max(\text{worst 30-day average}, 0) + 3\cdot\max(\sigma_{\text{paid}}, \sigma/3))$. It then takes a Wang (2000) expectation with $\lambda = 0.5$ and adds a 15% annual cost of capital on the limit over the 30-day term τ:

$$P = (E^{Wang}_{\lambda}[X] + r_c \tau L) / (1 + r_c \tau) \tag{2}$$

3. Data

For each venue we read the cumulative borrow index at 2,848 Ethereum blocks spaced 6,500 blocks apart (about one day), from block 7,930,000 (June 10, 2019) to block 26,435,500 (November 2026). Each read goes through one Multicall at that block against a public archive node. The venues and readings are: Compound v2 cUSDC (`borrowIndex`, updated to the block with `borrowRatePerBlock`, and `totalBorrows`); Aave v2 and Aave v3 (`getReserveNormalizedVariableDebt`, and the total supply of the variable- and, for v2, stable-debt tokens). We interpolate the log index to each UTC midnight. The daily realized rate is then $\exp(365 \cdot \Delta \ln I) - 1$, which is exactly what a borrower paid that day, compounding included. This avoids sampling the instantaneous rate, which can move many times within a day.

The **spliced index** follows the venue that held the most USDC borrowing: Compound v2 until May 21, 2021, Aave v2 until November 13, 2023, and Aave v3 after that. Aave v3 USDC is the product's actual settlement source. Volume is the sum of outstanding USDC borrowing across all three venues.

Table 1. Daily realized USDC borrow APY, percent.

Series	From	Days	Mean	Median	SD	P99	Max day	Max 30-day	ρ(1)	ρ(30)	Avg borrow
Compound v2	Jun 2019	2,716	7.54	4.59	9.17	32.5	128.3	65.2	0.97	0.27	$493m
Aave v2	Dec 2020	2,176	9.87	7.92	9.94	58.7	84.1	64.9	0.95	0.51	$724m
Aave v3	Jan 2023	1,389	6.29	5.25	3.13	17.1	22.8	15.1	0.94	0.64	$1696m
Spliced index	Jun 2019	2,716	6.33	5.14	3.75	17.2	31.5	15.6	0.95	0.66	$1940m

ρ(k) is the autocorrelation of daily rates at lag k days. Compound and Aave v2 extremes mostly come after each venue lost its lead and its USDC market became small or frozen. The spliced index avoids those periods.

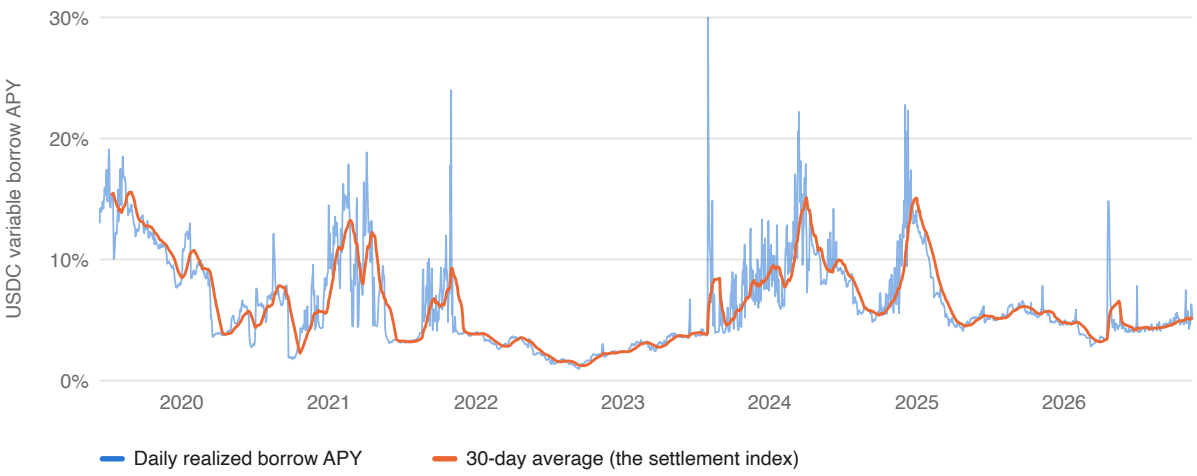

Figure 1. Spliced USDC variable borrow APY: daily realized rate and its 30-day average, which is the quantity the cover settles on. The daily axis is clipped at 30%.

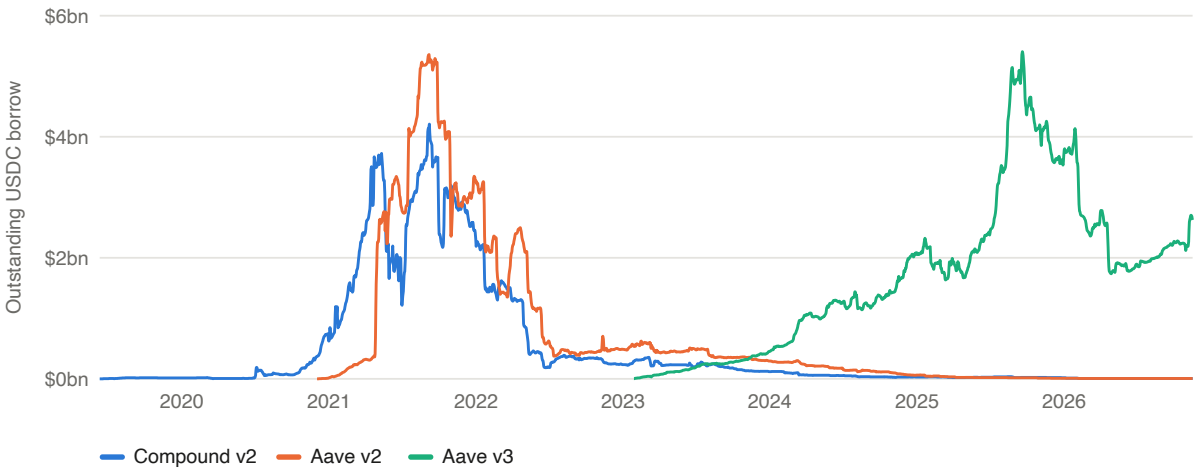

Figure 2. Outstanding USDC borrowing by venue. Borrowing peaked near $5 billion in 2021, fell to below $1 billion in 2023, and has since rebuilt on Aave v3.

4. The distribution of the index

The spliced daily rate has a mean of 6.33% and a median of 5.14%, with a right tail to 31.5% (July 31, 2023). Averaging over 30 days removes most of the spikes: the rolling 30-day average never went above 15.6% (Figure 3). The series is highly persistent (ρ(30) = 0.66, ρ(90) = 0.37). The lending curve's design produces this: rates stay pinned near the kink until utilization breaks through, then stay high until new supply arrives.

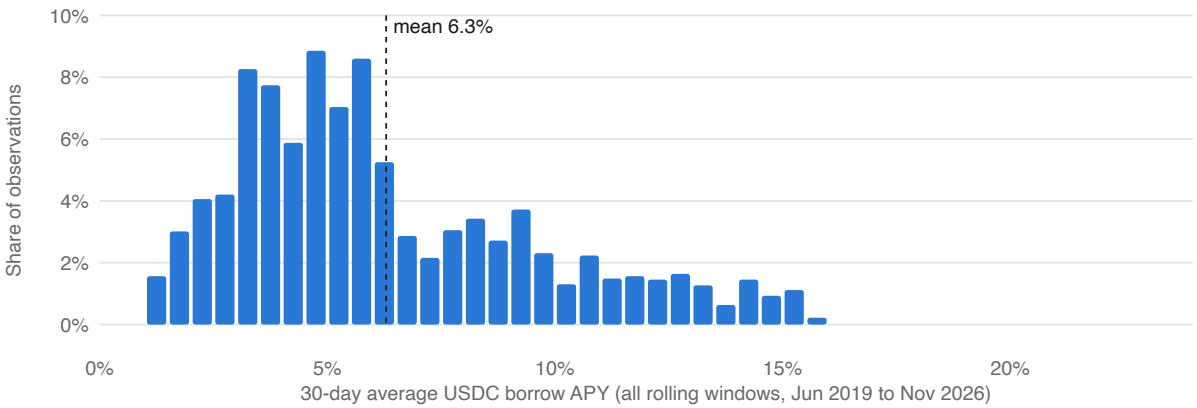

Figure 3. Distribution of all rolling 30-day averages of the spliced index. The mass sits between 2% and 8%, with a thin right tail to about 15.6%.

Table 2. The five worst non-overlapping 30-day periods of the spliced index.

Period	30-day average	Worst day	Context
Jul 29, 2019 – Aug 27, 2019	15.58%	18.5%	Early Compound, thin USDC market
Jun 14, 2019 – Jul 13, 2019	15.49%	19.1%	Early Compound, thin USDC market
Mar 2, 2024 – Mar 31, 2024	15.12%	22.2%	March 2024 rally, leveraged long demand
Dec 1, 2024 – Dec 30, 2024	15.08%	22.8%	December 2024 rally
Jan 23, 2021 – Feb 21, 2021	13.24%	17.9%	Early 2021 bull market

5. Backtest design

Every day t from June 10, 2020 (one year after the sample starts) to October 18, 2026 (the last day whose 30 days are fully observed), we sell a cohort of 30-day cover. The cohort is priced only on the 365 days before t, using the production code unchanged: `capLevels` and `quoteRateCap` from `protect-pricing.ts`. Each cohort covers θ/30 of outstanding USDC borrowing on day t, so cover in force follows θ times outstanding borrowing. We set θ = 5% as the base case. All dollar results scale linearly in θ. This gives 2,322 cohorts and an average of $112m in force against $2238m of outstanding borrowing. A cohort's realized payout comes from the 30 days after its sale, which the pricing never saw.

6. Results

Table 3. Underwriting results by cap, θ = 5% of outstanding USDC borrowing, June 2020 – October 2026.

Cap	Avg cap	Premium/yr	Claims/yr	Loss ratio	Rate on line	Cohorts paid	Avg capital	Return on capital	Worst drawdown
K0	5.94%	$3.47m	$0.55m	15.9%	25.9%	44%	$0.63m	460%	$239k
K1 (default)	6.98%	$2.72m	$0.24m	8.9%	20.4%	19%	$0.70m	357%	$142k
K2	8.02%	$2.11m	$0.12m	5.6%	15.8%	10%	$0.75m	267%	$104k
K3	9.07%	$1.62m	$0.07m	4.6%	12.1%	7%	$0.79m	196%	$89k

Rate on line is premium divided by limit. Return on capital is the annual underwriting result divided by average capital reserved (limit minus premium, summed over cover in force). It ignores the yield on idle reserves.

At the default cap, 19% of cohorts paid something, but payments were small compared with the premium: 81% of cohorts paid nothing, and Figure 4 shows that 84% of paying cohorts used less than a third of the limit. The loss ratio is 8.9%, and the 90% block-bootstrap interval (90-day blocks, 2,000 resamples) is 5%–15%. Even at the top of that interval, claims are well below premium. Reserved capital averaged $0.70m, or 0.62% of cover in force. That ratio is very small because the layer is thin: one month of a 10-point layer is less than 1% of notional.

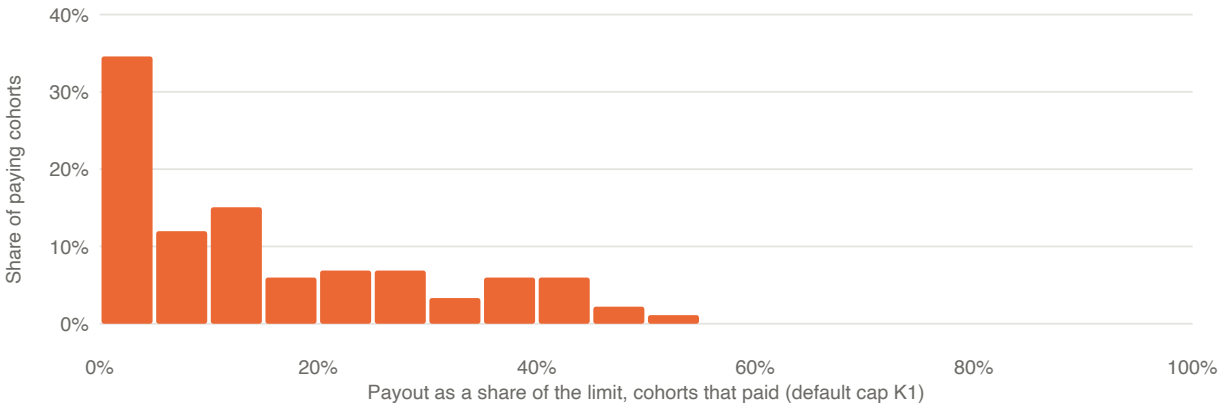

Figure 4. Payouts as a share of the limit, for default-cap cohorts that paid anything (19% of all cohorts).

Table 4. Default cap by calendar year of sale, θ = 5%.

Year	Avg USDC borrow	Cover in force	Realized avg rate	Premiums	Claims	Loss ratio	Avg capital
2020 (205 d)	$184m	$9m	6.00%	$0.13m	$0.07m	56%	$0.06m
2021	$5031m	$252m	6.70%	$7.27m	$0.52m	7%	$1.47m
2022	$1895m	$95m	2.36%	$2.73m	$0.00m	0%	$0.55m
2023	$836m	$42m	5.31%	$0.19m	$0.18m	96%	$0.33m
2024	$1411m	$71m	9.53%	$1.42m	$0.53m	37%	$0.46m
2025	$3068m	$153m	5.75%	$4.51m	$0.00m	0%	$0.89m
2026 (291 d)	$2367m	$118m	4.53%	$1.10m	$0.24m	22%	$0.86m
Total				**$17.33m**	**$1.54m**	**9%**	

The worst year was 2023 (loss ratio 96%). The trailing year was 2022's bear market, so caps were set low and cheap, and rates then rose through them in mid- and late 2023. 2021 wrote the most premium because borrowing was largest. 2022 and 2025 paid nothing.

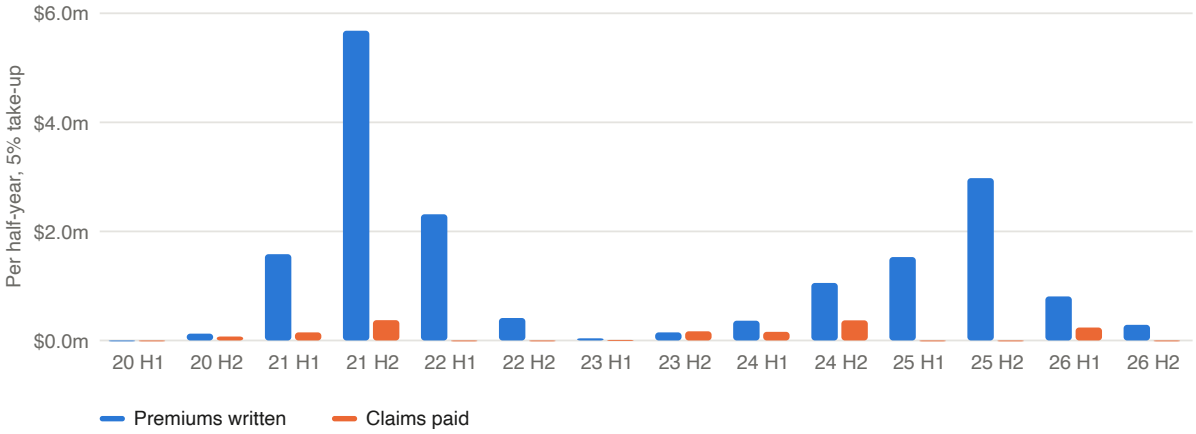

Figure 5. Premiums written and claims paid per half-year at the default cap, θ = 5%.

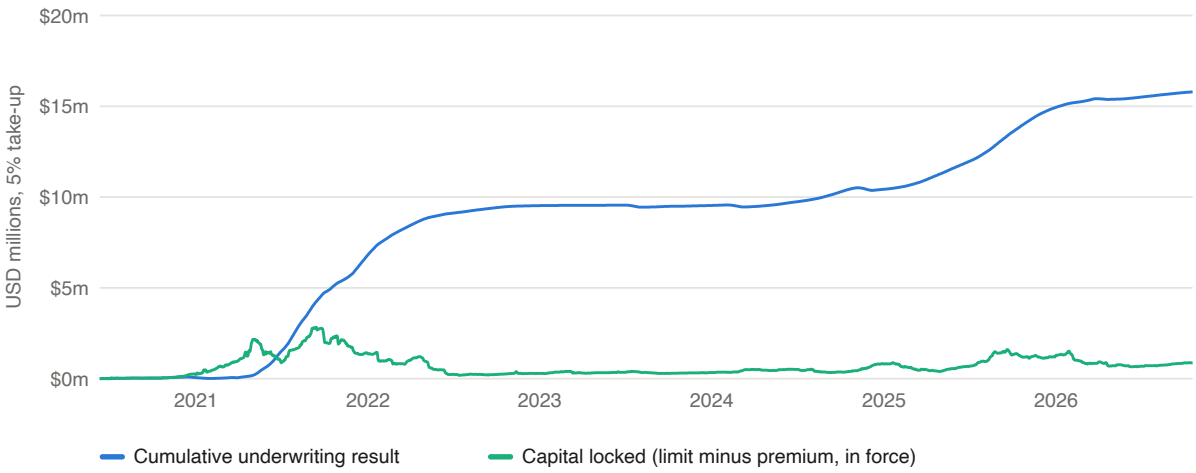

Figure 6. Cumulative underwriting result and reserved capital at the default cap, θ = 5%. The largest peak-to-trough fall in the cumulative result was $142k.

7. Why the price is high

Averaged over all default-cap cohorts, the premium was 20.4% of the limit, the trailing-year burn cost (the plain average payout over the 365-day window) was 8.6%, and realized payout was 3.0%. The gap therefore has two parts of similar size:

> **Loading:** premium / burn cost = 2.37×. This comes from the stress month (weight 1/24 on a payout near the full limit, worth up to 4.2% of the limit before distortion), the Wang distortion, and the capital charge (1.2% of the limit for 30 days).
>
> **Forecast bias:** burn cost / realized = 2.88×. The trailing-year estimate overstated forward claims.

We call the second effect a **spike hangover** (Figure 7). A rate spike enters the trailing window and stays there for 365 days, so it inflates the burn cost for a full year. Meanwhile caps are re-an-

chored to a current 30-day average that has already come back down, so spike-era windows clear the new, lower caps easily and count as large historical payouts. But spikes cluster and seldom repeat within the year. The burn cost therefore stays high long after the conditions that produced it have passed. Persistence makes this worse, not better: $\rho(30) \approx 0.66$ means the next 30 days look like the last 30, not like the trailing year.

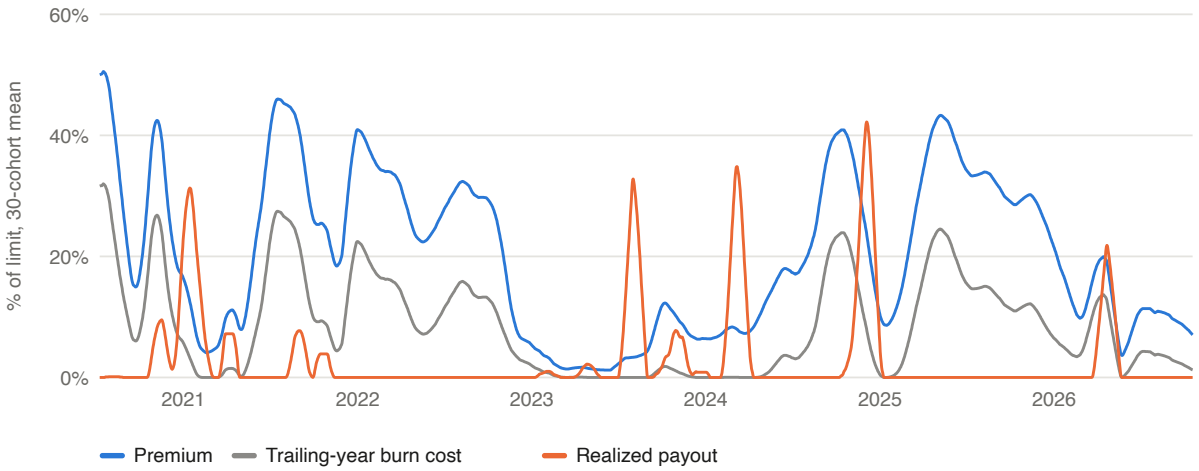

Figure 7. Premium, trailing-year burn cost and realized payout as shares of the limit, default cap, 30-cohort means. Burn cost jumps after each spike and stays raised for a year, while realized payouts return to zero.

7.1 Alternative calibrations

Table 5 re-runs the backtest with lighter loadings and with a *conditional* burn cost. The conditional version applies each historical 30-day-ahead change to today's level instead of using historical levels. That is a simple way to respect persistence.

Table 5. Default cap under alternative pricing parameters, $\theta = 5\%$. Claims are identical across rows; only premiums change.

Variant	λ	Stress weight	Cost of capital	Burn cost	Premium/yr	Loss ratio	Rate on line	Return on capital
Production	0.5	1/24	15%	levels	$2.72m	8.9%	20.4%	357%
Lighter	0.25	1/48	15%	levels	$1.90m	12.7%	14.2%	217%
Light	0.1	1/96	10%	levels	$1.49m	16.2%	11.1%	157%
Conditional	0.5	1/24	15%	changes	$1.69m	14.3%	13.0%	186%
Conditional, lighter	0.25	1/48	15%	changes	$0.94m	25.6%	7.3%	83%

No parameter set gets the loss ratio above about 26%. Lighter loadings bring the price down by up to half, and conditioning on changes removes part of the hangover, but the remaining forecast bias is structural. A longer lookback, using all history to date, made things worse in our tests: the loss ratio fell to 4.2%, because 2021-era spikes then stay in the window permanently.

7.2 Pricing frontier

Table 6 scales every premium by a common factor, which is what an experience-rated adjustment on top of the engine would do. The factor is fitted in sample, so this is a guide to the range, not a validated price.

Table 6. Default cap with all premiums multiplied by a factor, θ = 5%.

Price factor	Rate on line	Premium/yr	Loss ratio	Return on capital	Worst drawdown	Drawdown in years of premium	Losing years
1.00×	20.4%	$2.72m	9%	357%	$142k	0.05	0 of 7
0.50×	10.2%	$1.36m	18%	139%	$235k	0.17	2 of 7
0.35×	7.1%	$0.95m	25%	85%	$266k	0.28	3 of 7
0.25×	5.1%	$0.68m	35%	51%	$322k	0.47	3 of 7
0.15×	3.1%	$0.41m	59%	19%	$475k	1.16	4 of 7
0.10×	2.0%	$0.27m	89%	3%	$553k	2.03	4 of 7

The frontier is steep. At 0.25× the loss ratio is 35%, the pool still earns 51% a year on reserved capital, and it loses money in 3 of 7 calendar years. Getting the loss ratio into the 55–65% band typical of commercial insurance takes a factor of about 0.15×. At that level the pool earns 19% on capital but loses in 4 of 7 years, and its worst drawdown is 1.2 years of premium. Full collateral makes this a question of earnings volatility, not solvency. The choice depends on how much year-to-year variance underwriters will accept, and on the fact that the stress month has to stay in the price for regimes the sample never saw.

8. The buyer's view

A buyer who rolls the default cap every 30 days on $1 of borrowing pays the floating rate, minus payouts, plus premiums. Table 7 compares that with the unhedged cost over the same cohorts, in percent of notional per 30 days.

Table 7. 30-day borrowing cost per $1 of notional, with and without default-cap cover, production pricing.

	Mean	Standard deviation	99th percentile	Worst cohort
Unhedged	0.474%	0.244%	1.184%	1.243%
With cover	0.616%	0.219%	1.110%	1.208%

Net of the payouts it received, the cover raised the average 30-day cost by 14 bp of notional, about 1.74% a year. In exchange it trimmed the worst cohort from 1.24% to 1.21% and the 99th percentile by 7 bp. On this index the 30-day averaging already does most of the smoothing a buyer wants, and with the cap anchored near the current level, what the cover protects against is a regime shift within one month. Those were rare and modest in 2020–2026. A borrower who mainly fears a 20%+ month can buy the higher cap K_3 for 12.1% of the limit instead of 20.4%, though in this sample it paid in only 7% of cohorts.

9. Market size and capital

Revenue and capital are linear in take-up. Using average outstanding USDC borrowing of $2238m over the sample:

Table 8. Annual premiums, claims and reserved capital at the default cap, by take-up.

Take-up θ	Cover in force	Premium/yr (production)	Premium/yr (0.25×)	Claims/yr	Capital (production)	Capital (0.25×)
1%	$22m	$0.54m	$0.14m	$0.05m	$0.14m	$0.17m
5%	$112m	$2.72m	$0.68m	$0.24m	$0.70m	$0.86m
10%	$224m	$5.45m	$1.36m	$0.48m	$1.39m	$1.73m
25%	$559m	$13.62m	$3.41m	$1.21m	$3.48m	$4.32m

For context, the whole $2238m of USDC borrowing paid about $142m of interest a year at the sample's average rate. Insuring all of it at production prices would cost $54m a year, about 38% of the interest bill. Both figures are upper bounds on demand, and real take-up would be far lower.

10. Limitations

Index definition. The product settles on the mean of daily APYs from a public data aggregator for Aave v3. We use index-implied realized rates, which match it closely on Aave v3 but are not identical. Before November 2023 the spliced index is a proxy built from other venues. **Few independent tail events.** The sample holds perhaps five independent spike episodes. The bootstrap interval reflects this, but heavier regimes than any seen are possible. **Constant take-up.** Real demand would rise after spikes, which is exactly when trailing prices are highest, and fall in calm periods. We do not model this or the adverse selection it implies. **Reflexivity.** A large hedging market could change borrower behavior near the kink. **One asset, one chain.** USDT, DAI and other chains are not covered.

11. Conclusion

Over six years of out-of-sample cohorts, the production engine priced USDC rate caps safely: the pool was never short of capital and lost money in no year. But it priced them about 11× above realized claims. Roughly half the gap is deliberate loading. The other half is a trailing-window artifact that keeps every spike in the price for a year. A buyer would have paid about 2.4% of notional a year for limited reduction in variability. Our recommendation is to keep the engine's structure (empirical burn cost, explicit stress, distortion, capital charge) and correct its level. That means a conditional burn cost, a lighter distortion, and an experience-rated multiplier that starts near 0.25–0.35× and is credibility-weighted toward live claims as they accumulate (Bühlmann and Gisler, 2005). The data and code to reproduce every table are in the accompanying scripts.

References

1. Aave (2020). *Aave Protocol Whitepaper V2.0*.
2. Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3(1–2), 167–179.
3. Bühlmann, H. and Gisler, A. (2005). *A Course in Credibility Theory and its Applications*. Springer.
4. Gudgeon, L., Werner, S. M., Perez, D. and Knottenbelt, W. J. (2020). DeFi protocols for loanable funds: Interest rates, liquidity and market efficiency. *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 92–112.

5. Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17(3), 1217–1241.
6. Leshner, R. and Hayes, G. (2019). *Compound: The Money Market Protocol*. Whitepaper.
7. Wang, S. S. (2000). A class of distortion operators for pricing financial and insurance risks. *Journal of Risk and Insurance*, 67(1), 15–36.

Reproduction: `chain.mjs` (archive reads), `prep_lending.mjs` (daily index), `bt.ts` (backtest, imports the production `protect-pricing.ts`), `analysis.mjs` and `paper1.mjs`. Figures are computed, not illustrative. Nothing here is investment advice.