CoW DAO

Value Distribution Mechanism

Research Report (V2)

Prepared for: CoW DAO Core Team

Prepared by: Aragon Research

Date: April, 2026

1. Executive Summary

This report evaluates value distribution mechanisms for the COW token and recommends a shortlist of options to be further evaluated into the quantitative modelling part of the engagement. The two models are the current buyback-and-hold as baseline as well as a revenue-funded buyback-and-distribute model.

We derived evaluation criteria by combining Aragon's industry experience with four years of CoW DAO forum sentiment rather than imposing them externally. The community has consistently demanded revenue-funded rewards, visibility to holders, simplicity of participation, and solver-incentive neutrality. These constraints disqualify voting escrow and emissions-dependent mechanisms narrowing the viable design space to buyback-based models.

We applied these criteria across 11 protocols spanning five mechanism types: buyback/burn, buyback/distribute, buyback/hold, voting escrow, and social escrow. Three findings shaped the shortlist. First, the two closest structural analogues to CoW, 1inch and Paraswap, are clear failure cases. Both attempted escrow-style staking on DEX aggregator margins; 1inch produced an 8% stake rate with rewards funded by resolver discretion, and Paraswap's ultimately abandoned the mechanism. Second, buybacks alone cannot overcome emissions (Jupiter spent $70M against $1.2B in unlocks) and buyback-and-hold generates less visibility of the value accrual even at scale (Aave's $1M/week program coincided with ETH underperformance). Third, distribute models show empirical evidence of value retention: Sky and GMX achieved ~70% restake rates on distributed rewards, meaning most yield compounds into stake rather than hitting the market.

CoW's distributable margin is real but modest. The modelling phase runs a weekly Monte Carlo simulation over a two-year horizon, stress-testing both policies across stochastic COW price, revenue, and solver-cost paths. It reports achievable staking APR, circulating supply trajectories, and sell-through under varying distribution ratios, participation rates, and revenue assumptions.

2. Methodology

This research was conducted in four phases, each producing inputs for the next: understanding CoW's internal economics and political constraints, building a structured evaluation framework, cross-protocol empirical analysis, and mechanism shortlisting for quantitative modelling.

Note on constraints: The available dataset is both small and heterogeneous. The 11 protocols reviewed operate across different market regimes, revenue scales, and mechanism lifetimes. No two implemented the same mechanism under comparable conditions. We cannot isolate the causal effect of any single mechanism with statistical confidence. The methodology of this research is adapted to this constraint, instead of forcing rigour where the domain does not support it. Directional claims are grounded in observed patterns, stated with appropriate uncertainty, and accompanied by the context needed to assess their transferability to CoW.

Phase 1: Protocol & Sentiment Analysis. We mapped COW token value flows across solver compensation, fee-to-buyback pipeline, and net COW flows post-CIP-74 (introduction of volume fee). In parallel, we reviewed CoW DAO forum discussions from 2022 to 2026, tracking recurring themes and evolving positions on value accrual.

Rather than imposing external evaluation criteria, we derived them inductively from revealed community preferences. This produced the constraint set that governs the rest of the analysis.

Phase 2: Decision Framework. Combining sentiment-derived constraints and our experience in mechanism design, we defined evaluation criteria spanning qualitative dimensions (legibility, visibility, solver-incentive compatibility) and observable outcomes (effective supply reduction, value accrual efficiency, market beta outperformance) as detailed in Section 5. Given the highlighted dataset constraints, the framework structures judgment rather than replaces it.

Phase 3: Cross-Protocol Benchmarking. We analysed 11 protocols including direct DEX aggregation competitors and covering the full mechanism taxonomy (buyback-burn, buyback-distribute, buyback-hold, escrow). We could not examine clean examples of real yield distribution as Sky discontinued USDS distributions in favor of in-kind. For each, we collected observable outcome data: stake rates, buyback volumes, supply dynamics, price performance relative to ETH/SOL. A consolidated comparison table is presented in Section 6, accompanied by a proximity analysis assessing each

benchmark's transferability to CoW across business model, revenue scale, and token distribution dimensions.

Phase 4: Shortlisting & Modelling. We shortlisted finalist mechanisms for deeper quantitative modelling against CoW's revenue profile, free cash generation scale, and solver-neutrality constraint. The modelling phase is delivered in a separate report document. We are not constructing a protocol P&L forecast but rather basing the distribution for simulation purposes on historical results.

3. CoW Protocol: Current State

Revenue

CoW Protocol collects fees in the surplus token across all supported chains through three mechanisms:

- Surplus fee on out-of-market limit orders (50% of surplus, capped at 0.98% of volume) since CIP-34
- Quote improvement fee on market orders (50% of positive price improvement, same cap) since CIP-61
- Volume fee of 2bps on all trades since CIP-74, later reduced to 0.3bps for correlated pairs

The first two mechanisms are conditional on surplus and quote improvement, meaning some batches generate zero protocol fees. The volume fee guarantees a revenue floor per batch, improving revenue predictability.

Cost Structure

Solver compensation is the primary cost line. Solvers are paid weekly in COW for three functions:

- A fixed per quote reward regardless of which solver settles
- An execution reward determined by a second-price auction, per batch capped at a percentage protocol fees generated
- A consistency reward, based on user participation on the total number of executed orders.

Gas reimbursements and payments to integration partners are the other significant variable costs, although they are passthrough given the volume based fee structure. Partners charge their own fees on

attributed orders. CoW collects these gross and remits the partner's share, retaining a service fee of 25% by default, but can be negotiated. While OpEx is funded separately from the Core Treasury it should be considered when looking at the Protocol's net position. CIP-79 determined the 2026 development budget, CIP-82 determined the Grants program size and CIP-70 (Ops Safe) and CIP-64 (Foundation mandate) determined the remaining Foundation related OpEx.

The Buyback Loop

The Treasury Team targets buying back 1.2x weekly solver distributions to maintain a net reduction in circulating COW. This gets executed via weekly TWAP orders on mainnet. This is an operational policy, not a protocol-level rule and holds as long as fee revenue is sufficient relative to solver costs at prevailing COW prices.

Net Position

The protocol moved from negative gross margins in 2023 through breakeven in mid-2024 to consistently positive margins by late 2025. Team vesting completed February 2026, removing a large source of structural sell pressure. A new team allocation was approved under CIP-83 streams 50M COW, roughly 1.04M COW per month - with further milestone based incentives - as well as 6M COW in grants. Treasury yield on stables provides additional income.

While the funds might flow through separate loops, structurally the revenue flow looks as follows: Income after gas and partner costs funds solver rewards, excess COW to the buyback multisig and residual cash flow to the Managed Treasury. Residual cash ideally exceeds OpEx budget. The distributable margin available for any value accrual mechanism is real but modest - the mechanism we recommend must optimize for signal efficiency per dollar.

4. Forum Sentiment Analysis

We reviewed CoW DAO forum discussions from 2022 to 2026, supplemented by Snapshot voting patterns, tracking how community opinion on value accrual has evolved. Six observed key themes are summarised below:

- **Buybacks are invisible.** While the buyback program has been running for over a year and has achieved net negative COW supply impact the market hardly knows about its existence. The token needs a legible value story, not because the underlying economics are broken, but because the signal is not reaching holders or prospective buyers. This is core tension to resolve.
- **Grow the pie first, distribute later.** Through 2024, the dominant position was that revenue should fund growth rather than flow to tokenholders. While recently the stance has weakened, it remains the consensus.
- **No inflationary staking rewards.** Ignited by Cobie's ApeCoin essay in 2022 and echoed by Hasu in the CoW forums, there is strong consensus that staking rewards funded by token inflation are destructive and attract mercenary capital while diluting long-term holders. Any staking mechanism must be funded by protocol revenue. There is positive sentiment around amplified voting rights as an additional utility for staked COW.
- **VE complexity is a cost.** Since 2022, community members have consistently argued that voting-escrow systems are unnecessarily complex for CoW's current state. While the "too early" objection no longer holds (CoW is now one of the largest DEX aggregators by volume) the complexity objection remains valid.

Note: This sentiment has been validated by observed outcomes at 1inch and Paraswap, where VE-adjacent systems in competitors produced low participation and minimal value capture.

- **Do not sell COW to fund operations.** In early 2024, there was sharp criticism of selling COW from the treasury to fund protocol operations while simultaneously running a buyback program. Many argued that doing buybacks and then selling a similar amount of COW from the treasury is contradictory. The clear preference is to fund operations from non-COW assets (ETH, stables) and preserve COW for value accrual purposes only.
- **Solver economics are non-negotiable.** With CIP-74 and surrounding discussions in late 2025, it became evident that solver profitability is a real and active concern. Any distribution mechanism that diverts revenue from solver compensation will be detrimental to CoW-Solvers alignment and consequentially to the success of CoW as a leading dex aggregator. It became clear that distribution should be funded from the DAO's share of revenue, *after* the solver rewards, ensuring

a solver-incentive neutrality guarantee. This principle extends to any mechanism that might increase cost to solvers (i.e. opportunity cost of locked capital).

Taken together, these themes define a narrow design space: revenue-funded, visible to holders, simple to understand and participate in, non-dilutive and operationally neutral to solvers. The decision framework in Section 5 translates these constraints into evaluation criteria.

5. Decision Framework

Based on the constraints identified in Section 4 and our experience in mechanism design, we derived the following mechanism analysis dimensions. Assessment is summarized in Section 6

1. **Legibility.** How easily can a participant understand the mechanism? The market performance of a mechanism is contingent on a low context, low effort participant understanding how it drives value. A simple buyback and burn is easier to understand than a Vote-escrow-based MetaDEX fork.
2. **Visibility.** How clearly can holders see value accruing to them? Distinct from legibility - a staking APR is very visible whereas a buyback and hold can often be missed.
3. **Sustainable Value Source.** Is the mechanism funded by protocol revenue, token inflation, or a combination? Per the community's hard constraint, inflationary rewards are disqualified.
4. **Effective Supply Reduction.** How effectively is the mechanism in reducing circulating or available supply? Burns are permanent, locks are temporary with expiry risk, staking without lockups relies on participation friction alone. Measured via observed changes such as circulating supply, stake rate, and buyback-to-emission ratio.
5. **Value Accrual Efficiency.** Per dollar spent, how much value is retained by holders versus leaked back to the market? Burn generates no yield signal. Hold preserves optionality but delivers little immediate impact. Distribute concentrates on stakers but may be resold. Distributing with lockup concentrates further and reduces sell pressure. This forces comparison on capital efficiency, not just mechanism type.
6. **Implementation Ease.** Measured by smart contract and auxiliary software complexity.
7. **Solver-Incentive Compatibility.** If implemented for CoW Swap, would the mechanism preserve, improve, or degrade solver economics? Worsening the solver experience with an onerous value

accrual mechanism, or not utilising the DAO's share of revenue, may have the second-order effects that solvers may exit and make the protocol less profitable.

8. **Market Beta Outperformance.** Is there evidence of the mechanism causing outperformance vs. ETH/SOL beyond what macro conditions explain? The hardest criterion to assess given dataset constraints.

6. Mechanism Evaluation

A summary of each reviewed protocol against the decision criteria is available here. The data sources used for the presented analysis are consolidated on the following page

6.1 Lock + Delegation (1inch)

Data Source

VE-style staking with resolver delegation. Stakers lock 1INCH (1 month to 2 years) to receive Unicorn Power (UP), a decaying token delegated to resolvers in exchange for rewards.

Rewards are funded by resolvers, who distribute 1INCH from their own revenue to delegators, supplemented by a Foundation token allocation. Protocol-level fees do not flow to stakers. A mid-2025 proposal to route trading fees to the DAO failed over concerns about degrading resolver economics. Staking participation averaged 18% of supply. Staker APY averaged 8% (range: 5-10%). Underperformed ETH.

- **Legibility** is low because participation requires a multi-step process with UP decay mechanics.
- **Visibility** is low because the reward source is opaque and there's no protocol-level aggregate APY.
- **Sustainable value source** is weak: rewards come from resolver 1INCH balances and a Foundation token allocation, not protocol fees.
- **Effective supply reduction** is moderate at 18% average participation.
- **Value accrual efficiency** is low: 8% average APY for stakers on a protocol processing billions in volume, with no protocol revenue in the mix.
- **Implementation ease** is poor due to UP decay math, early exit penalties, and resolver delegation infrastructure.

- **Solver-incentive compatibility** - introduces a lot of complexity as staker revenue flows through solvers
- **Market beta** underperformed ETH over the mechanism's lifetime.

Relevance to CoW. Closest structural analogue. Despite reasonable participation, single-digit APY and opaque funding structurally cap the mechanism's appeal. Models where yield depends on solver discretion are fragile.

6.2 Social Escrow (Paraswap)

Data Source

Activity-based staking with market maker delegation. Stakers lock PSP to receive sePSP and earn rewards for protocol-positive actions (trading, market making, referrals). Paid in ETH from protocol fees. Token fell 99.9% from ATH and was abandoned for VLR rebrand in September 2025.

Funding was hybrid: protocol fees (distributed as ETH) plus PSP emissions (6% of supply allocated to market maker rewards). The revenue component was insufficient to sustain meaningful yields alone, and emissions dwarfed fee revenue. Two staking flavors existed - sePSP1 (pure PSP) and sePSP2 (80/20 PSP/ETH Balancer BPT, boosted rewards). Unstaking required a 28-day cooldown. Staking participation averaged 21.5% with average APY of ~16%, reflecting thin per-participant distributions during low-volume periods. Supply grew from 194M to 2B (10x) between Jan 2022 and Sep 2025 as emissions outpaced any impact staking had. Underperformed ETH.

- **Legibility** is low due to bespoke, multi-action reward logic and two staking variants.
- **Visibility** is moderate: yields were visible but thin per participant.
- **Sustainable value source** is medium, combining ETH-denominated protocol fees and PSP emissions but emissions dwarfed fees. Interestingly the only protocol distributing ETH in addition to native token.
- **Effective supply reduction** was negligible as supply grew 10x despite reasonable staking participation.
- **Value accrual efficiency** is low. Despite decent participation, yield was mostly emissions driven.
- **Implementation ease** is poor given the activity tracking and boost multiplier infrastructure required.
- **Solver-incentive compatibility** - low impact, especially if staked COW can be used for the bond

- **Market beta** underperformed ETH.

Relevance to CoW. Closest structural analogue alongside 1inch. Activity-based reward design doesn't solve the underlying problem: if protocol revenue is insufficient and the mechanism is propped up by emissions, participation metrics look healthy while the token still collapses.

6.3 Buyback (& burn later) (Jupiter)

Data Source

Revenue-funded buybacks split between immediate burn and the Litterbox Trust - a discretionary reserve holding bought-back tokens for potential future use. Staker rewards exist via Active Staking Rewards (ASR) which are quarterly airdrops funded from a community token allocation, not from the buybacks themselves. Net-Zero Emissions proposal passed February 2026 with 75% approval, ending further emissions.

Cumulative buyback-and-burn reached approximately 220M JUP; ASR distributed approximately 200M JUP over the same period. Over the buyback program's first year, approximately $70M was spent against approximately $1.2B in unlocks hitting the market. While the team is not happy with the results of the program, it is hard to assess whether a different mechanism would have been more successful given the ratio of inflation to free cash. A one-time 3B JUP burn in January 2025 (vesting cuts and reserve reductions) removed 30% of total supply. A further 135M Litterbox burn executed November 2025.

- **Legibility** is medium. Buyback is simple, but the burn/Litterbox split adds ambiguity.
- **Visibility** is medium. Buybacks are on-chain but ASR distributions discrete and hard to express as APR
- **Sustainable value source** is revenue-funded at 50% of protocol fees.
- **Effective supply reduction** is low as buybacks covered only ~6% of new supply during the period.
- **Value accrual efficiency** is low. With $70M in buybacks against $1.2B in unlocks, mechanism choices have little to no impact.
- **Implementation ease** is low complexity.
- **Solver-incentive compatibility** - mechanism has no impact on solver incentives.
- **Market beta** underperformed SOL.

Relevance to CoW. Jupiter is operating in a similar vertical as a DEX aggregator, however the revenue scale is materially different (~ $280M in 2025). The primary finding is that buybacks as a signal get netted out by large unlock schedules which is a problem CoW doesn't have.

6.4 Buyback (and Burn) (Hyperliquid)

Data Source

Hyperliquid is running an automated buyback via the Assistance Fund, programmatically funded by 97% of protocol fees. It is the strongest buyback program in crypto by absolute spend. Average buyback is ~70k HYPE per day (~2.1M/month). Staking exists separately and is funded from a 100M HYPE emission pool over 10 years, not from buybacks. Circulating supply reduced by ~14% since November 2024. Monthly unlocks of 9.9M HYPE begin in 2026, which will test the buyback's supply-reduction capacity against a 4.7x larger emission rate.

- **Legibility** is high: protocol fees auto-accumulate in the Assistance Fund, auto-buy HYPE, tokens burned. No governance decisions required.
- **Visibility** is medium-high: the burn is narrative-powerful and on-chain trackable (~8-10% annualized burn rate). Staking yield is funded through inflation on a separate mechanism, unrelated to the accrual one.
- **Sustainable value source** is revenue-funded directly tying the mechanism to the success of the underlying business.
- **Effective supply reduction** is strong to date (-14% since launch), but 2026 unlocks will exceed buyback absorption (9.9M/month unlocks vs 2.1M/month buybacks) and will stress-test the trajectory.
- **Value accrual efficiency** is mechanically low as a burn mechanism doesn't directly relay value to tokenholders but rather increases their relative share of the supply.
- **Implementation ease** is low. The burn pattern is easily replicable.
- **Solver-incentive compatibility** - mechanism has no impact on solver incentives.
- **Market beta** outperformed ETH significantly. Buybacks don't drive price mechanically, but large-scale automated mechanisms create conditions for speculative capital to price the asset as structurally scarce.

Relevance to CoW. Hyperliquid has fundamentally different revenue model, user base, and competitive dynamics. With revenue scale approximately 80x CoW's, the specific mechanism (97% of fees to

automated burn) is not replicable. The transferable signal is that programmatic, governance-free mechanisms provide credibility in a way discretionary ones don't, and that scale relative to float creates narrative conditions that compound the direct economic effect.

6.5 Buyback + Distribute to Stakers (Sky)

Data Source

Automated buyback of SKY on the open market, distributed to stakers in SKY. The mechanism has been live since MakerDAO's rebrand to Sky in August 2024 and succeeded earlier designs that used auction burns. The buybacks are funded by the protocol retained earnings (called "surplus"). Stake rate averages 60% with an average APY around 14.5%. Notably 72% of distributed rewards are restaked, although a large element of this may be due to the ability to borrow/loop USDS against staked SKY collateral. Cumulative buybacks reached ~1.8B SKY (~$112M) over the past year. In March 2026, the protocol reduced buyback rates to strengthen stablecoin reserves amid macro pressure, demonstrating the discretionary nature of the allocation.

- **Legibility is high** from a holder perspective: stake SKY, receive SKY yield funded by protocol surplus. Stakers can borrow USDS against their stake, adding utility without complicating the reward mechanic.
- **Visibility is high.** Displayable APR, continuous distributions, on-chain trackable buyback program.
- **Sustainable value source is revenue-funded** from protocol surplus (excess above the Surplus Buffer target), though allocation is governance-discretionary.
- **Effective supply reduction is meaningful but conditional.** 60% average stake rate with 72% restake rate indicates most distributed yield is recycled into stake rather than sold, despite being paid in the volatile token. Staking has no duration lock so staked supply is unlikely to be considered "out of circulation"
- **Value accrual efficiency is medium.** Yield is paid in SKY, which exposes the mechanism to the reflexive dynamic the COW community has flagged as a concern. However the high restake rate offsets this.
- **Implementation ease is moderate.** Mechanism is tractable but non-trivial.
- **Solver-incentive compatibility** - low impact, especially if staked COW can be used for the bond
- **Market beta outperformed ETH** over the mechanism's lifetime.

Relevance to CoW. Closest match among the benchmarks to the design space the community has signalled: revenue-funded, visible APR, no lockups, no inflation. The iteration path is relevant as earlier Sky mechanisms (flap burn, SBE LP) were abandoned for reasons that map to CoW's stated constraints (procyclicality, illegibility, weak holder signal). Sky converged on the current design by eliminating those failure modes. A constraint to consider is that Sky's surplus is sizable, and that some elements of the mechanism are unique to Sky's position as a stablecoin issuer.

6.6 Discretionary Buyback + Distribute to Stakers (Ether.fi)

Data Source

Multisig-executed buybacks of ETHFI distributed to sETHFI stakers. It is funded by 5% of protocol revenue and 100% of ETH unstaking fees and supplemented by a $50M buyback program funded from treasury that triggers when ETHFI trades below $3. Average buyback size is ~1.5M ETHFI/month against ~4.5M ETHFI/month in vesting unlocks. Stake rate averages 6.6%, APR averages 14%. Unstaking requires up to a 10-day withdrawal queue.

- **Legibility is medium.** Stake ETHFI, receive ETHFI from buybacks. Simple core mechanic, but fully discretionary timing and sizing, with no programmatic commitment or predictable schedule.
- **Visibility is high.** Staking with visible yield from a defined revenue source is a clear, immediate signal to participants.
- **Sustainable value source is revenue-funded** but structurally insufficient. Buybacks cover roughly one-third of monthly vesting unlocks.
- **Effective supply reduction is low.** 6.6% average stake rate and 4.5M/month unlocks against 1.5M/month buybacks. Staking spiked when withdrawal fees were added as a funding source and then normalized.
- **Value accrual efficiency is medium.** Staking demand responded more to feature announcements than buyback indicating the mechanism's value is in the signal, not the mechanical supply reduction.
- **Implementation ease is low complexity.** Manual multisig transactions.
- **Solver-incentive compatibility** medium as the 10-day withdrawal queue creates friction for solvers.
- **Market beta** underperformed ETH.

Relevance to CoW. Some relevant findings: First, discretionary, unpredictable buyback execution weakens the mechanism's signal value compared to programmatic approaches like Sky's. Second, the conditional price-triggered buyback ($50M below $3) is a capital efficiency pattern potentially worth considering for CoW. Lastly the 10-day withdrawal queue adds friction without VE complexity, but also without observable impact.

6.7 Voting Escrow (Aerodrome)

Data Source

VE-locking with fee distribution to voters. Users lock AERO for up to 4 years to receive veAERO, vote weekly to direct emissions to liquidity pools, and earn 100% of trading fees accruing to staked Aero LPs from voted pools. Lock rate is around 50%. Revenue exceeded emissions briefly (September 2025); for most of the protocol's life, emissions outpaced fees.

- **Legibility is low.** Multiple interacting elements - lock, vote, direct emissions, earn fees, manage decay
- **Visibility is high.** APR on locked positions is clearly displayed and drives participation.
- **Sustainable value source is emissions-dominant.** Trading fees are real but emissions outpaced them for most of the protocol's history. The Aero Fed (post-epoch 67) gives voters control over emission rates but doesn't resolve the structural imbalance.
- **Effective supply reduction is mixed.** ~50% lock rate, but locks don't recapture all emissions. First 4-year lockups are yet to expire in August 2027.
- **Value accrual efficiency is high but reflexive.** The flywheel compounds in both directions. 80% token price drawdown despite stable TVL demonstrates the downside.
- **Implementation ease is low** This is one of the highest complexity reviewed. Gauge system, voting, boost mechanics, lock decay, emissions scheduling, voter-controlled monetary policy.
- **Solver-incentive compatibility would be poor.** Complex capital lockup creates opportunity cost.
- **Market beta.** Outperformed ETH slightly but with high volatility.

Relevance to CoW. Deprioritised on three grounds: community has rejected VE complexity since 2022, the flywheel depends on emissions CoW's community has explicitly disqualified, and capital lockup requirements could violate solver-incentive neutrality. The transferable observation is narrow: voter-directed emissions create strong engagement mechanics, but VE is structurally incompatible with CoW's constraints.

6.8 Buyback and Distribute (GMX)

Data Source

Revenue-funded buyback of GMX distributed to stakers using 27% of protocol fees. Average buyback ~120K GMX/month against ~49K/month in emissions. Stake rate averages 76%, but it's good to note that rate was already in that range with the earlier esGMX reward system, so buyback contribution was marginal. In March 2026, GMX acknowledged "limited effectiveness" and paused staking reward distributions.

- **Legibility is high.** Protocol fees buy GMX, distribute to stakers. Permissionless execution via FeeHandler allowing anyone to trigger the buyback onchain.
- **Visibility is high.** Visible yield of distributions to stakers, replaced by buyback rate as primary metric.
- **Sustainable value source** funded by 27% of protocol revenue from trading, liquidation, & borrowing fees.
- **Effective supply reduction is high.** Buybacks exceeded emissions 2.4:1, and 70% of distributed rewards were restaked. Net sell-through was ~30% of distributions.
- **Value accrual efficiency is medium.** While sellthrough of rewards was given as reason for discontinuing the distributions, data shows 70% restake rate. Value leakage could be explained by other factors (July 2025 hack, low buyback to total float, overall downward market).
- **Implementation ease is medium.** Oracle-dependent FeeHandler with permissionless execution.
- **Solver-incentive compatibility** - low impact, especially if staked COW can be used for the bond
- **Market beta underperformed ETH** over the mechanism's lifetime.

Relevance to CoW. GMX's headline 76% stake rate is not a function of the design. It was inherited from the previous staking system. The more useful finding is that the distribute mechanism achieved a 70% restake rate despite continuous price decline. For CoW, this strengthens the case that a distribute model can retain most of its value in stake, provided the APR clears the opportunity cost threshold.

6.9 Buyback (and Hold) (Aave)

Data Source

AFC-managed buybacks at $1M/week ($50M/yr budget, proposed reduction to $30M/yr after ~25% borrow fee decline). Roughly 1.3% of supply bought back over a year and sent to Ecosystem Reserve with no defined distribution timeline. Staking rate is around 20% but the staking mechanism is unrelated to the buybacks.

- **Legibility is medium.** The buyback action is simple (multisig purchases AAVE from secondary markets), but the program sits inside significant governance infrastructure. What happens to bought tokens is undefined.
- **Visibility is low.** Tokens accumulate in the Ecosystem Reserve. The same reserve funds Umbrella staking rewards but the relationship between buyback accumulation and holder benefit is opaque.
- **Sustainable value source is** revenue-funded with conservative guardrails.
- **Effective supply reduction is minimal.** 1.3% of supply acquired in a year. The bought tokens are not burned and could re-enter circulation at governance's discretion.
- **Value accrual efficiency is medium.** No immediate value to holders. Optionality for future distribution, burns, or strategic use is preserved but unrealised.
- **Implementation ease is low technical complexity** (multisig) but meaningful governance overhead.
- **Solver-incentive compatibility -** no impact.
- **Market beta.** Underperformed ETH slightly with high volatility despite continuous $1M/week buyback.

Relevance to CoW. Aave is the direct structural comparison to CoW's existing buyback-and-hold program. The core finding: buyback-and-hold at moderate scale preserves optionality but generates no holder-visible signal.

6.10 Quasi-buyback and Burn (Uniswap)

Data Source

Permissionless, oracle-less burn mechanism. Protocol fees accumulate in the TokenJar contract; the only way to withdraw fees is to permanently burn UNI via the Firepit contract. Live since December 2025 across v2, v3, and Unichain. A one-time burn of 100M UNI (~$596M, 10% of supply) was executed from the treasury alongside fee activation. Annualised ongoing burn is estimated at ~4.4M UNI (~0.4% of

supply). It is overshadowed by a 20M UNI annual growth budget. Too early for definitive performance assessment.

- **Legibility is medium.** The core concept (burn UNI to claim fees) is simple but the distinction from protocol-driven buybacks is easy to miss, and the fee capture structure varies across sources
- **Visibility is low.** No individual holder yield. The burn is on-chain trackable but the self-adjusting rate (accelerates when UNI price falls, decelerates when it rises) is not visible to casual observers.
- **Sustainable value source is revenue-funded** across diversified fee sources (v2, v3, v4 pending, Unichain).
- **Effective supply reduction is early but positive.** Ongoing burns at ~4.4M/year are modest relative to the 20M/year growth budget, making the net ongoing effect inflationary
- **Value accrual efficiency is medium.** Permanent supply removal, but no yield signal.
- **Implementation ease is medium.** Smart contracts are simple but the mechanism design (oracle-less, decentralised execution, tiered fee capture across protocol versions) is architecturally sophisticated.
- **Solver-incentive compatibility is not applicable.** Uniswap has no solver equivalent.
- **Market beta** is too early to assess

Relevance to CoW. Limited direct transferability. Uniswap explicitly chose burn over distribution for regulatory and tax reasons. It is too early to assess the performance of the program. The self-adjusting burn rate (accelerates when price falls, decelerates when it rises) is a useful capital efficiency concept that connects to the Lido NEST discussion on purchase timing optimisation.

6.11 NEST Buyback (Lido)

Proposed automated buyback deploying protocol revenue into an LDO/wstETH Uniswap v2-style LP position held by the DAO. Activates only when ETH exceeds $3,000 and annualised revenue exceeds $40M, with a $10M annual cap. Neither condition currently met. A separate $20M one-off discretionary buyback began executing April 2026 via CoW Swap and centralised exchanges. Not yet operational as an automated mechanism.

- **Legibility is medium.** Intent is clear (automated buybacks when affordable). Activation thresholds are well-defined. The LP deployment mechanic adds complexity beyond a simple buyback.

- **Visibility is high.** Clear program with on-chain auditability. Caveat is that the system is not yet active.
- **Sustainable value source is revenue-funded** with built-in self-throttling. Lido's 2025 revenue declined 23% YoY to $40.5M, barely at the threshold, demonstrating the guardrails' relevance.
- **Effective supply reduction is medium.** Tokens deployed as protocol-owned LP are practically removed from circulation but likely to still be counted by dashboards as "circulating supply".
- **Value accrual efficiency is medium.** LP deployment builds protocol-owned liquidity alongside the buyback, addressing a real problem (LDO on-chain depth is ~$90K at ±2%).
- **Implementation ease is medium.** Technically complete (NEST v1 developed). Parameter calibration requires financial engineering.
- **Solver-incentive compatibility** - no impact
- **Market beta** is too early to assess

Relevance to CoW. Primarily a purchase-mechanics reference. Conditional activation thresholds that self-throttle when the protocol can't afford the buyback are directly applicable given CoW's variable revenue and fixed opex. The LP co-investment model addresses token depth issues CoW also faces.

7. Deprioritisation Rationale & Shortlist

Deprioritised Mechanisms

Voting Escrow (Aerodrome/Curve). Deprioritised on three grounds. First, the community has consistently signalled that VE complexity is too high. Second, VE systems require capital lockup that creates direct opportunity cost for solvers. Third, VE systems depend on continuous emissions to sustain their flywheels. CoW's community has explicitly rejected inflationary rewards, and CoW's gross margin cannot sustain the required emission-to-fee ratios.

Aggregator Escrow (1inch/Paraswap). The closest structural analogues to CoW and clear failure cases: rewards funded by resolver balances or emissions rather than protocol revenue, resolver economics turned net negative at 1inch, and Paraswap's mechanism was abandoned after a 99.9% token decline. The DEX aggregator model operates on margins too thin to generate a VE-style flywheel without degrading solver/resolver competitiveness. However, the delegation-to-solver mechanic could deserves further exploration as a standalone design element.

Buyback & Burn (Hyperliquid). While the program has delivered impressive results it is non-replicable. CoW is already allocating its distributable margin to buybacks at 1.2x solver emissions. The question is not allocation but signal: at CoW's scale, a burn produces no holder-visible yield, no displayable APR, and no participation mechanic. From our perspective, the separation of holding tokens versus burning tokens is tough to isolate. Even Hyperliquid's "burn" is contingent on no hardforks of the network unlocking the burned HYPE in the future.

Shortlisted for Modelling

After applying the full criteria set against CoW's specific constraints (modest gross margin, solver-incentive neutrality, community preference for visibility and simplicity), we recommend modelling a combination of the following:

1. A baseline Buyback and Hold model, under variable market conditions and budgets.
2. A buyback and distribution model, alongside a simple staking system (in both COW and real yield)

The core comparison: does distribution to stakers generate sufficient visible yield to justify the increased circulating supply versus holding in treasury? Empirical restake rates from comparable mechanisms suggest a large share of distributed rewards are recycled into stake rather than sold, though whether CoW's achievable APR sustains comparable retention is a modelling question. The modelling phase will test achievable APR at CoW's revenue level across stake rate and growth assumptions.

Two enhancement layers will be explored alongside the base mechanism:

- Staking mechanics that reduce the remaining sell-through without punitive lockups.
- Purchase optimisation via market-timed buybacks (Lido NEST model), improving effective yield per dollar of gross margin deployed. CoW Swap is the natural execution venue.