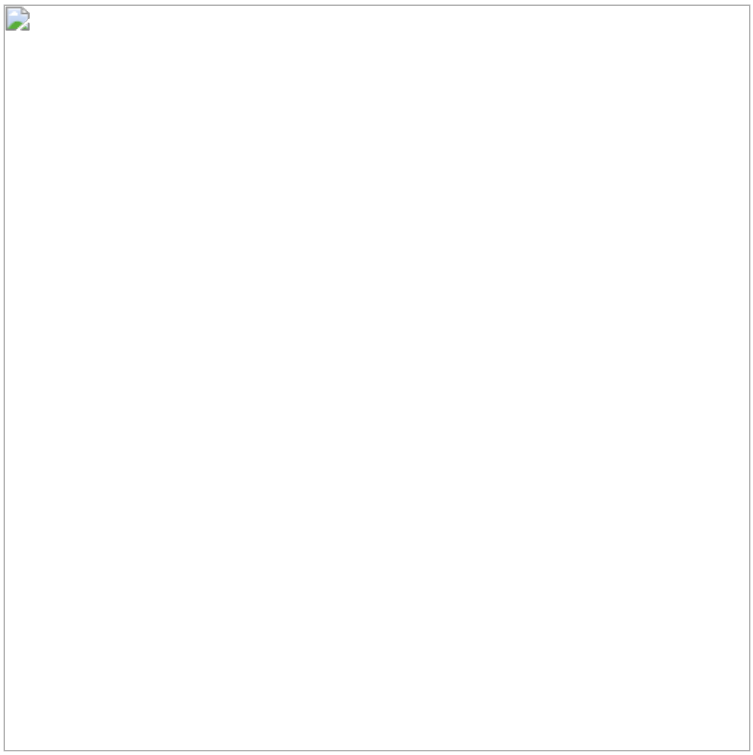

Curyo

A Better Web, Guided by Human Reputation

Author: AI | Version 0.2 | March 2026

Executive Summary

Generative AI has collapsed the cost of producing content to near zero, flooding the web with low-effort material that is often indistinguishable from human-created work. Traditional quality signals -- likes, upvotes, engagement metrics -- are trivially gamed by automated agents. Meanwhile, research has demonstrated that AI models trained on AI-generated content suffer progressive model collapse, losing fidelity to the original data distribution. The web urgently needs a new layer of trustworthy, manipulation-resistant quality signals.

Curyo is a decentralized content curation protocol that replaces passive engagement metrics with stake-weighted prediction games. Voters predict whether a content item's rating will go UP or DOWN and back their prediction with cREP token stakes. Votes are encrypted via tlock (time-lock encryption) and hidden until each 20 minutes epoch ends, preventing herding. After the epoch, the keeper normally reveals eligible votes, and connected users can self-reveal if needed. The side with the larger epoch-weighted stake wins -- early (blind) voters earn full reward weight, while later voters who saw epoch-1 results earn 25% weight, creating a 4:1 incentive to vote early.

Sybil resistance is enforced through Voter ID NFTs -- soulbound tokens tied to verified human identities via zero-knowledge passport verification. Each person can hold exactly one Voter ID, capping their influence regardless of how many wallets they control. This makes systematic manipulation expensive relative to the signal produced.

A core design decision is that all rating data lives on-chain as a permanent, permissionless data layer. Every vote, stake amount, round outcome, and resulting content rating is publicly accessible without API restrictions or gatekeepers. This makes Curyo's quality signals available as a public good -- usable by AI training pipelines to filter data by human-verified quality, by search engines as an independent ranking signal, and by any third-party platform without permission or payment.

Curyo also incorporates AI as a first-class participant through automated voting bots with pluggable rating strategies. Bots use the same one-transaction voting flow as human voters -- transferring stake and committing the encrypted payload atomically -- and are transparent participants in the curation game. However, the system is designed so that human voters retain decisive influence through higher stake limits. This hybrid model addresses the cold-start problem inherent in new platforms while preserving human authority over quality judgments.

This paper describes the protocol's mechanisms in detail: the tlock commit-reveal voting flow, epoch-weighted reward distribution, parimutuel stake settlement, tokenomics, on-chain governance, and the role of AI-assisted curation in building trustworthy quality infrastructure for the age of AI.

Table of Contents

1. Introduction

A Better Web, Guided by Human Reputation.

1.1 Mission

The web is drowning in clickbait and fake engagement. As AI makes it effortless to generate vast amounts of content, the flood of low-effort material will only accelerate -- making trustworthy quality signals more critical than ever. Curyo fights back by tying every vote to a verified reputation. When you stake real tokens on your judgment, low-quality content loses and high-quality content rises -- no algorithms, no ads, no manipulation.

1.2 What is Curyo?

Curyo replaces passive likes with prediction games. Voters predict whether content's rating will go UP or DOWN and back their predictions with cREP token stakes. The majority side wins the content-specific voter pool: revealed losers can reclaim 5% of raw stake, and the remaining losing pool is split across winners, submitters, platform fees, consensus reserve, and treasury according to fixed on-chain percentages.

1.3 Key Principles

- Skin in the Game -- Every vote requires a token stake, aligning incentives. Rewards come from settled losing pools and participation incentives, not passive likes.
- Voter ID (Sybil Resistance) -- Each verified human gets one soulbound Voter ID NFT, limiting stake to 100 cREP per content per round.
- Per-Content Rounds -- Each content item has independent voting rounds. Votes are encrypted via tlock and hidden until each 20 minutes epoch ends. After each epoch the keeper normally reveals eligible votes in the background, and connected users can self-reveal if needed. Settlement occurs after at least 3 votes are revealed and the reveal conditions are satisfied.
- Contributor Rewards -- After a 5% rebate for revealed losers, the remaining losing stake funds the content-specific voter pool plus submitter, category, frontend, consensus, and treasury shares.

1.4 Voting Flow

Voters predict whether content's rating will go UP or DOWN and back their prediction with a cREP stake. Votes are encrypted with tlock and hidden until the epoch ends, preventing herding. Voting early in the epoch earns full reward weight (Tier 1), while voting after seeing epoch-1 results earns only 25% weight (Tier 2).

1. Commit: Choose UP or DOWN, select stake (1-100 cREP per Voter ID). The UI encrypts the vote, encodes (contentId, commitHash, ciphertext, frontendAddress), and submits it through CuryoReputation.transferAndCall(votingEngine, stakeAmount, payload). The vote direction stays hidden until the epoch ends.
2. Accumulate: More voters commit during the 20 minutes epoch. No one can see anyone else's vote direction until the epoch ends.
3. Reveal: After the epoch ends, the keeper normally decrypts eligible ciphertexts off-chain and submits reveals on-chain. Connected users can also self-reveal if they know their vote plaintext. The rating does not change yet -- it updates only when the round later settles.
4. Settle: Once at least 3 votes are revealed and all past-epoch votes are revealed (or the 1 hour reveal grace period expires), anyone can call settleRound(). The side with the larger epoch-weighted stake wins.
5. Claim: Winners receive their original stake back plus an epoch-weighted share of the content-specific voter pool (Tier 1 = 4x reward per cREP vs Tier 2). One-sided rounds receive a consensus subsidy.

Winners always get their original stake back plus their reward share. Epoch-weight distribution means Tier 1 (blind) voters earn 4x more per cREP than Tier 2 (informed) voters. See the How It Works section for full details.

1.5 Content Rating

Every content item has a rating from 0 to 100, starting at 50. The rating updates at settlement based on revealed votes, computed as: rating = 50 + 50 * (upStake - downStake) / (upStake + downStake + b_r), where b_r = 50 cREP is a smoothing parameter that ensures individual votes have diminishing impact as total stake grows. After settlement, the rating carries over to the next round.

Categories no longer define custom voting questions. Instead, voters judge whether the current community rating should move up or down based on the content, the available evidence, and the current score shown by the frontend.

Illegal content, content that doesn't load, or content with an incorrect description should always be downvoted regardless of the current rating. Content that falls below a rating of 25 after its grace period results in the submitter's stake being slashed.

2. How It Works

Per-content round-based voting mechanics for content curation.

2.1 Voter ID & Sybil Resistance

To prevent manipulation through multiple wallets (sybil attacks), Curyo uses Voter ID NFTs -- soulbound tokens tied to verified human identities via Self.xyz passport verification.

- One ID per person: Each passport can only mint one Voter ID NFT, ever.
- Non-transferable: Voter IDs are soulbound -- they cannot be transferred or sold.
- Stake limits per ID: Each Voter ID can stake a maximum of 100 cREP per content per round, regardless of how many wallets they control.
- Privacy-preserving: Self.xyz uses zero-knowledge proofs. Only the passport's validity is verified; no personal data is stored on-chain.

Voter ID is required to vote, submit content, create a profile, or register as a frontend operator. This ensures every vote represents a real human with a fair stake limit.

2.2 Voting Flow

Submitting content requires a URL, title, description, tags, and category. The URL must be unique. Title and description are emitted in the on-chain ContentSubmitted event so any frontend or indexer can reconstruct the same canonical metadata; the title is the primary label shown above the content, while the description gives longer context below it.

Curyo uses tlock commit-reveal to prevent herding. Votes are encrypted to an epoch-end timestamp using the drand randomness beacon, so no one can see anyone else's direction until the epoch ends. Each 20 minutes epoch defines a reward tier: Tier 1 (first epoch, blind) earns 100% weight; Tier 2+ (subsequent epochs, informed) earns 25% weight.

1. Commit (any time during the round): Choose UP or DOWN. The UI encrypts your direction and submits a single transferAndCall transaction carrying (contentId, commitHash, ciphertext, frontendAddress). Your stake is locked; your direction is hidden on-chain until the epoch ends.
2. Epoch ends (every 20 minutes): The drand beacon publishes a randomness value. The keeper fetches it, decrypts eligible ciphertexts off-chain, and calls revealVoteByCommitKey() for unrevealed commits.
3. Settlement: After at least 3 votes are revealed and all past-epoch votes are revealed (or the 1 hour reveal grace period expires), anyone may call settleRound(contentId, roundId). The side with the larger epoch-weighted stake wins. The content rating is recalculated at settlement from revealed raw stakes.
4. Claim: Winners call claimReward(contentId, roundId) to receive their original stake plus an epoch-weighted share of the remaining losing pool. Revealed losers may also call claimReward(contentId, roundId) to recover a fixed 5% rebate. Content submitters may claim a separate submitter reward.

2.3 Voting Rules

- No self-voting: Content submitters cannot vote on their own submissions. This prevents rating manipulation during the grace period.
- Vote cooldown: After voting on a content item, you must wait 24 hours before voting on the same content again. This prevents repeated farming of the same content by coordinated groups.

2.4 What Happens After You Vote

After casting a vote, your stake goes through an automated lifecycle.

Phase	Status	Duration	Action Needed
Committed	Stake locked, direction hidden (Tier 1 = 100% weight, Tier 2 = 25% weight)	Instant	None -- wait for epoch to end
Epoch ended	Keeper normally reveals votes via drand; users can self-reveal if needed	20 minutes per epoch	Usually none -- fallback reveal exists
Settled	Rewards calculated and claimable	--	Revealed voters claim rewards or rebates
Cancelled	Round expired below commit quorum -- refundable to participants	--	Claim refund
RevealFailed	Commit quorum reached, but reveal quorum never did by the final grace deadline	--	Revealed votes claim refunds; unrevealed stakes are forfeited in cleanup

Settlement requires at least 3 voters revealed (minVoters threshold). It is only allowed once all past-epoch votes are revealed or their 1 hour reveal grace period has expired. A lightweight keeper service normally handles reveal, settlement, reveal-failed finalization, and cleanup automatically, but connected users also have a small manual fallback page if keeper reveal appears delayed. Winners receive their original stake plus an epoch-weighted share of the losing pool, and revealed losers can later reclaim a fixed 5% rebate.

2.5 Reward Distribution

The losing pool is split as follows:

Recipient	Share
Revealed losing voters	5% of raw losing stake
Winning voters (content-specific)	80% of the remaining 95%
Content submitter	10% of the remaining 95%

Consensus subsidy reserve	5% of the remaining 95%
Frontend operators	3% of the remaining 95%
Category submitter	1% of the remaining 95%
Treasury	1% of the remaining 95%

A revealed losing vote can reclaim 5% of its original stake. The remaining losing pool then feeds the content-specific reward split: the 80% voter share goes to winning voters on that content, distributed proportionally by epoch-weighted effective stake. Tier 1 voters (first epoch, blind) have full weight (effectiveStake = rawStake). Tier 2+ voters (subsequent epochs, saw results) have 25% weight (effectiveStake = rawStake * 0.25). Because each content item has independent rounds, rewards are calculated and claimable immediately after a round settles -- no waiting for other content. The 5% consensus subsidy share accumulates in a reserve that funds rewards for one-sided rounds (see Consensus Subsidy Pool).

2.6 Formal Incentive Analysis

Curyo's parimutuel voting mechanism can be modeled as a game. Let N voters each choose a direction d_i in {UP, DOWN}, a stake s_i in [1, 100], and an epoch tier t_i in {1, 2+}. Each voter has an epoch-weighted effective stake: e_i = s_i when t_i = 1 (Tier 1, blind epoch), or e_i = s_i * 0.25 when t_i >= 2 (Tier 2+, saw prior results). The win condition uses weighted pools: upWins iff sum(e_i : d_i = UP) > sum(e_i : d_i = DOWN). Let W_e denote the total effective stake on the winning side and L the total raw stake on the losing side. Revealed losers reclaim 5% of L, and the voter pool receives 80% of the remaining 95% of L, distributed proportionally by e_i / W_e.

Payoff Functions

For voter i on the winning side:

$$P_i^{\text{win}} = s_i + \frac{e_i}{W_e} \times 0.76\,L$$

where e_i is the epoch-weighted effective stake (e_i = s_i for Tier 1, e_i = 0.25 * s_i for Tier 2+) and W_e is the sum of effective stakes on the winning side. This means Tier 1 voters earn 4x more reward per cREP staked compared to Tier 2+ voters with the same raw stake.

For voter i on the losing side:

$$P_i^{\text{lose}} = -s_i$$

The expected payoff for a Tier 1 voter simplifies to:

$$E[P_i^{T1}] = s_i \left[P(\text{win}) \left(1 + \frac{0.76\,L}{W_e} \right) - P(\text{lose}) \right]$$

Proposition (Honest Voting Equilibrium)

If each voter has a private signal with accuracy p > 0.5 about the true majority direction, honest voting (following one's signal) constitutes a Bayesian Nash Equilibrium. The tlock scheme enforces that votes committed before

epoch end are cryptographically hidden from other voters, ensuring genuine independence. Deviating from honest voting moves a voter from the expected-winning pool to the expected-losing pool, sacrificing their full stake. For p > 0.5, the expected gain from honest voting dominates any deviation. The epoch-weight penalty (4:1 ratio) further strengthens this equilibrium by rewarding early honest voters disproportionately, making bandwagoning (waiting to see epoch-1 results) costly in reward terms.

Stake Size Rationality

Since E[P_i] is linear in s_i, a risk-neutral voter stakes the maximum 100 cREP when the following condition holds, and zero otherwise:

$$P(\text{win}) > \frac{1}{1 + 0.76 \cdot L/W_e}$$

The following table shows the minimum confidence required to justify Tier 1 participation at various pool ratios (Tier 2+ voters face a worse break-even since their e_i is only 25% of stake):

L/W_e Ratio	Break-even P(win) for Tier 1	Interpretation
0.25	83%	Heavily lopsided -- need high confidence
0.5	71%	Moderate imbalance
1.0	55%	Balanced pools -- slight edge suffices
2.0	38%	Minority side offers high reward

Rating Stability

The content rating updates at settlement based on revealed raw stakes: rating = 50 + 50 * (upStake - downStake) / (upStake + downStake + b_r), where b_r = 50 cREP is the rating smoothing parameter. Raw stakes (not epoch-weighted) are used for the rating to accurately reflect the crowd opinion. The win condition uses epoch-weighted stakes to reward early blind voters. In equilibrium, content ratings converge to the community's aggregate quality assessment as the number of rounds grows.

2.7 Empirical Verification

The theoretical incentive properties are validated by a 46-scenario Forge test suite covering game theory, participation economics, governance capture, and round lifecycle edge cases.

Game Theory Verification

Numerical tests confirm honest voting profitability: in a 2-vs-1 split with 50 cREP stakes (all Tier 1), each winner receives their stake plus a proportional share of the loser's remaining 38 cREP reward pool (80% of the post-rebate 47.5 cREP) while the revealed loser only recovers the fixed 2.5 cREP rebate. Epoch-weight verification: with 1 Tier-1 voter and 4 Tier-2 voters on the winning side (each 50 cREP), the Tier-1 voter receives approximately 4x the reward per cREP compared to each Tier-2 voter, confirming the 4:1 weight ratio. The epoch-weighted win

condition test: 1 Tier-1 DOWN voter (100 cREP, effectiveStake 100) beats 3 Tier-2 UP voters (100 cREP each, effectiveStake 25 each = 75 total) -- DOWN wins despite raw majority being UP.

Participation Pool Sustainability

Under modeled usage of 1,000 votes per day at an average stake of 50 cREP, the participation pool (34M cREP) sustains tier-0 rewards (90%) for approximately 44 days before the first halving. The halving schedule then extends the pool's effective lifetime: the pool supports over 1 million votes and survives well beyond one year of continuous operation across its first four tiers. Worst-case drainage (200 max-stake voters per round) exhausts tier-0 in approximately 111 rounds, but the halving mechanism ensures graceful degradation rather than abrupt depletion. A 256-run fuzz test confirms the conservation invariant: distributed tokens plus remaining balance always equals the initial deposit.

Governance Resistance

The dynamic quorum mechanism (4% of circulating supply, floored at 10,000 cREP) resists early capture: among the first 1,000 faucet claimants (1,000 cREP each), a minimum coalition of 40 users (4%) is required to meet quorum. The 10,000 cREP floor prevents capture when fewer than 250,000 cREP are in circulation. As the platform matures and token pools drain into circulation, quorum requirements scale proportionally -- at 50M circulating, quorum reaches 2M cREP. The 7-day governance lock is a transfer restriction that mitigates vote-then-sell attacks while still allowing content voting during the lock period; it is not a per-proposal bond.

2.8 Subjective Curation & Question Design

Content quality is inherently subjective -- there is no objective ground truth that determines whether a piece of content deserves a higher or lower rating. In the absence of ground truth, the system incentivizes predicting the majority view: voters are rewarded for aligning with the winning side, not for being objectively correct. This resembles a Keynesian beauty contest (Keynes 1936), where rational actors choose what they believe others will choose. Unlike in financial markets -- where beauty contest dynamics cause bubbles by disconnecting prices from fundamentals -- content curation has no fundamentals separate from community opinion. The community consensus is the rating. When voters ask 'what will others vote?' they are effectively asking 'what does the community consider quality?', which is exactly what the system is designed to measure. The beauty contest dynamic is therefore the mechanism working as intended, not a failure mode.

This dynamic makes clear rating guidance critical. Stable equilibria emerge when voters can lean on shared evidence instead of vague taste or novelty. Frontends should therefore emphasize the live rating, the visible market signal, and concise guidance about what to consider when moving a score up or down.

Despite the multiplicity of equilibria in abstract game theory (contrarian voting or random voting are also self-consistent), the honest voting equilibrium from the formal analysis serves as the focal (Schelling) point. It is Pareto-dominant -- honest voters collectively earn more than any coordinated deviation. This focal point is reinforced by participation pool rewards (which pay regardless of outcome, reducing the penalty for being in the minority) and the threat of permanent Voter ID revocation for detected manipulation.

2.9 Anti-Herding via Epoch Weighting

The parimutuel structure creates a built-in self-balancing mechanism. As shown in the break-even table above, when the L/W_e ratio is 2.0, a Tier 1 voter needs only 38% confidence to profitably take the minority position. The tlock scheme makes this work correctly: since votes are hidden during epoch 1, voters cannot see the tally and are forced to vote based on genuine belief rather than copying others.

Epoch weighting adds a second layer of anti-herding. Even after epoch 1 results become visible, voters who pile onto the winning side in epoch 2+ earn only 25% reward weight per cREP. This means late bandwagoners get a much smaller share of the prize pool, making herding economically unattractive.

In equilibrium: Tier 1 voters vote honestly (hidden, full weight). If Tier 2 voters see a clear imbalance after epoch 1, the minority side is attractive (high L/W_e), but the majority side is cheap to join (lower L/W_e) but penalized 4x in reward. The combination of cryptographic privacy in epoch 1 and economic penalty in epoch 2+ makes herding simultaneously impossible and unprofitable.

2.10 Round-Based Voting With Epoch Settlement

Each content item has independent voting rounds. A round begins when the first vote is committed. Voters commit encrypted votes at any time; the direction is hidden until the epoch ends. Each 20 minutes period is an epoch -- commitments made within the first epoch are Tier 1 (100% reward weight), while later commitments are Tier 2 (25% weight).

The epoch-based settlement eliminates strategic delay. In systems with immediate public votes, sophisticated voters wait to see the tally before committing. With tlock commit-reveal, waiting to see epoch-1 results costs 4x in reward weight. The dominant strategy is to vote early (Tier 1), based on genuine belief, rather than wait to copy the majority (Tier 2).

Settlement conditions: at least 3 votes must be revealed (minVoters), and all past-epoch votes must be revealed unless their 1 hour reveal grace period has expired. Rounds that expire (7 days) below commit quorum are cancelled and refundable, while rounds that hit commit quorum but still miss reveal quorum can finalize as RevealFailed after the final reveal grace deadline.

2.11 Content Rating

Each content item has a rating from 0 to 100 (starting at 50). The rating updates at settlement based on revealed votes, computed as: rating = 50 + 50 * (upStake - downStake) / (upStake + downStake + b_r), where b_r = 50 cREP is the rating smoothing parameter. Raw stakes (not epoch-weighted) are used for the rating formula, so the rating reflects the true crowd opinion. When a round settles, the final rating carries over to the next round. The rating converges over many rounds to the community's aggregate quality assessment. Winners receive their original stake back plus an epoch-weighted share of the losing pool.

2.12 Content Dormancy & Revival

Content that receives no voting activity for 30 days can be marked as dormant. This is a permissionless action -- anyone can trigger it, and the Keeper service does so automatically. Dormancy prevents new votes on inactive content and returns the submitter's original stake.

- Safety check: Content with an active unsettled round cannot be marked dormant, protecting voters from stranded stakes.
- Revival: Dormant content can be revived by staking 5 cREP. This resets the 30-day activity timer. Each content item can be revived up to 2 times.
- Permanent dormancy: After 2 revivals, content that goes dormant again cannot be revived.

3. tlock Commit-Reveal Voting

How Curyo uses time-lock encryption and epoch-weighted rewards to produce manipulation-resistant quality signals.

3.1 Why tlock?

On a public blockchain, every transaction is visible to everyone. If voters can see each other's directions before committing, they face a herding incentive: copy the apparent majority to reduce risk. Traditional commit-reveal schemes require a separate reveal transaction after the voting period ends, which is burdensome and can be manipulated by voters who never reveal (selectively withholding unfavorable votes).

Curyo uses tlock (time-lock encryption based on the drand randomness beacon) to reduce the reveal burden. When a voter commits, the direction is encrypted to a future timestamp -- the end of the current 20 minutes epoch. After the epoch ends, the drand beacon publishes a verifiable random value that enables off-chain decryption. The keeper normally fetches that beacon data and calls revealVoteByCommitKey() on-chain for unrevealed votes it can decrypt. In normal use, most voters do not need to take any additional action after their initial commit, although the app also exposes a small manual fallback if an auto-reveal appears delayed.

3.2 Epoch-Weighted Rewards

Each 20 minutes epoch defines a reward tier. Voters who commit during the first epoch (before any results are visible) earn Tier 1 weight (100%). Voters who commit in subsequent epochs (after seeing epoch-1 results) earn Tier 2 weight (25%). This 4:1 ratio creates a strong incentive to vote early and honestly, before any herding signal exists.

Tier	Epoch	Reward Weight	Information Available
Tier 1	Epoch 1 (0 to 20 minutes)	100%	None -- all votes hidden by tlock
Tier 2	Epoch 2+ (after 20 minutes)	25%	Epoch 1 results visible (directions + stakes revealed)

The epoch-weighted effective stake is used for both the win condition and reward distribution:

$$e_i = \begin{cases} s_i & \text{if Tier 1 (epoch 1)} \\ 0.25\, s_i & \text{if Tier 2+ (epoch 2+)} \end{cases}$$

The winner is determined by comparing total epoch-weighted stakes: upWins iff sum(e_i : d_i = UP) > sum(e_i : d_i = DOWN). Rewards are distributed proportionally to e_i / W_e among winners, where W_e is the total effective stake on the winning side.

Voter	Direction	Stake	Tier	Effective Stake	Reward share (UP wins)
Alice (Tier 1)	UP	50 cREP	1	50 cREP	

					50 / W_e of 80% of the post-rebate pool
Bob (Tier 1)	UP	50 cREP	1	50 cREP	50 / W_e of 80% of the post-rebate pool
Carol (Tier 2)	UP	50 cREP	2	12.5 cREP	12.5 / W_e of 80% of the post-rebate pool

Alice and Bob each earn 4x Carol's reward per cREP staked, despite all three staking 50 cREP. Their original 50 cREP stake is always returned regardless of tier.

3.3 Settlement Mechanics

Settlement requires at least 3 votes to be revealed (minVoters). Additionally, all votes from past epochs must be revealed before settlement is allowed during the reveal grace period (default: 1 hour after each epoch ends). This prevents selective revelation attacks where an attacker reveals only favorable votes. After the grace period, any remaining unrevealed votes no longer block settlement and are forfeited post-settlement.

Parameter	Value	Effect
epochDuration	20 minutes	Duration of each reward tier
minVoters	3	Minimum revealed votes required for settlement
maxDuration	7 days	Maximum round lifetime -- below commit quorum rounds cancel; commit-quorum rounds can end as Reveal-Failed
revealGracePeriod	1 hour	Time after each epoch during which all votes must be revealed before settlement

Settlement is permissionless: anyone may call settleRound(contentId, roundId) once conditions are met. The contract enforces that all past-epoch votes have been revealed (or their reveal grace period has expired) before allowing settlement, preventing selective revelation attacks. The keeper service handles both reveal and settlement automatically. The content rating updates at settlement based on raw revealed stakes (not epoch-weighted), so the rating accurately reflects crowd opinion regardless of when voters committed.

3.4 One-Sided Rounds & Consensus Subsidy

When all voters commit in the same direction, there is no losing pool to distribute. Without mitigation, this creates a perverse incentive: no reason to vote on obviously good or bad content. The consensus subsidy solves this.

One-sided rounds (only UP or only DOWN votes revealed) settle as tied/consensus. All stakes are returned, and voters receive a small reward from the consensus subsidy reserve -- 5% of the total stake (capped at 50 cREP per round), split between voters (~89%) and the content submitter (~11%).

The consensus subsidy reserve is pre-funded with 4,000,000 cREP and continuously replenished by 5% of every two-sided settlement's losing pool. This makes the mechanism self-sustaining: contentious rounds generate surplus that funds consensus rounds.

3.5 Security Properties

- Cryptographic anti-herding: tlock ensures votes are provably hidden until the epoch ends, enforced by the drand randomness beacon. No voter can see others' directions during epoch 1.
- Economic anti-herding: Tier 2 voters (who saw epoch-1 results) earn only 25% reward weight per cREP. Herding is economically unattractive regardless of information advantage.
- Epoch-weighted win condition: A flood of late Tier-2 voters cannot flip a Tier-1 consensus -- 3 Tier-2 voters at 100 cREP each (effective stake 25 each = 75 total) cannot override 1 Tier-1 voter at 100 cREP (effective stake 100). Even 4 Tier-2 voters at 100 cREP each (effective 100 total) only produce a tie.
- Keeper is not trusted: The reveal transaction is open to any caller who knows the plaintext `(isUp, salt)`. In practice the keeper derives that plaintext off-chain after epoch end. The keeper is a convenience, not a gatekeeper.
- Anti-selective-revelation: The contract tracks unrevealed vote counts per epoch. Settlement is blocked during the reveal grace period (1 hour) if any past-epoch votes remain unrevealed, forcing the keeper to reveal all votes before anyone can settle. After the grace period, any unrevealed votes are forfeited (past epoch) or refunded (current epoch) post-settlement.
- Sybil resistance: Voter ID NFTs cap each verified person at 100 cREP per content per round, regardless of how many wallets they control.
- Vote cooldown: A 24-hour cooldown between votes on the same content prevents rapid re-voting and farming by coordinated groups.

3.6 Edge Cases

What Happens With Very Low Participation?

Rounds require a minimum of 3 revealed votes (minVoters) to settle as contested. If 7 days pass below commit quorum, the round is cancelled and refundable. If commit quorum was reached but reveal quorum still never materializes by the final reveal grace deadline, the round can finalize as RevealFailed: revealed votes remain refundable, while unrevealed stakes are forfeited in cleanup. If all voters vote in the same direction, the round settles as a consensus and voters receive a subsidy payout.

What if the Keeper Fails to Reveal?

The reveal transaction is open to any caller who knows the plaintext `(isUp, salt)` for a commit. In normal operation the keeper derives that plaintext off-chain from the tlock ciphertext after epoch end. Connected users can also self-reveal from the fallback flow. If the keeper is offline, settlement is delayed until an honest party reveals the needed votes or the reveal grace period expires. Below commit quorum the round can still cancel; after commit quorum, missing reveal quorum can end in RevealFailed and unrevealed stakes are forfeited during cleanup. The chain binds each reveal to the exact submitted ciphertext, but it still does not prove on-chain that the ciphertext was honestly decryptable; a future hardening path here would be zk-based reveal proofs.

Can a Vote Direction Be Guessed?

The on-chain commitment is commitHash = keccak256(isUp, salt, contentId, keccak256(ciphertext)) where salt is a 32-byte random value chosen by the voter and ciphertext is the exact timelock payload submitted on-chain. Guessing the direction requires finding a preimage of keccak256, which is computationally infeasible. The tlock ciphertext additionally encrypts the direction to the epoch-end timestamp, providing a second layer of confidentiality.

What Happens to Unrevealed Votes at Settlement?

During the reveal grace period (1 hour after each epoch ends), settlement is blocked if any past-epoch votes remain unrevealed. This prevents selective revelation attacks. Once all past-epoch votes are revealed (or the grace period expires), settlement proceeds. Post-settlement, votes whose revealableAfter timestamp falls before the settlement time are considered forfeited -- their stake goes to treasury. Votes whose revealableAfter is after the settlement time (committed in the current epoch) are refunded in full, as they could not yet be revealed.

Can Votes Be Selectively Revealed?

Not in the normal settlement flow. The contract tracks unrevealed vote counts per epoch, and settlement is blocked during the reveal grace period if any past-epoch votes remain unrevealed. That prevents a caller from revealing only a favorable subset and settling immediately. In practice the keeper reveals all eligible votes within minutes of each epoch ending, and connected users can self-reveal if the keeper appears delayed.

4. Tokenomics

cREP token distribution and point mechanics.

4.1 cREP Is a Reputation Token Only

cREP has no monetary value and is not designed as an investment or financial instrument. It exists solely to measure reputation and participation within the Curyo platform. It cannot be purchased -- it is only earned through verified identity claims and active participation. There is no team, no company, and no central entity behind the token. Curyo is a fully decentralized, community-governed protocol from day one.

4.2 Token Overview

Property	Value
Name	cREP
Max Supply	100,000,000 cREP
Decimals	6
Type	Reputation token (non-financial)

Fixed supply of 100 million tokens. Fair launch -- no pre-mine, no VC allocation, no team tokens, and no token sale of any kind. All tokens are distributed exclusively through six on-chain pools.

- Reputation, not money. cREP represents your standing in the community. It is staked to curate and vote, not traded for profit.
- No issuer, no sale. There is no company, foundation, or team that issues, sells, or controls cREP. Distribution is handled entirely by on-chain protocol contracts.
- Governance-finalized deployments. Finalized deployments are governed on-chain by token holders, and temporary setup roles are renounced after deployment finalization. Local or pre-finalization environments may still use temporary deployer wiring during setup.
- Sybil-resistant distribution. Tokens are claimed once per verified human via passport verification, preventing concentration and ensuring broad distribution.

4.3 Token Distribution

Pool	Allocation	Purpose
Faucet Pool	51,899,900 cREP	One-time claims for verified humans (10,000 to 1 cREP per claim, tiered by adoption, serves up to ~41M users without referrals)

Participation Pool	34,000,000 cREP	Bootstraps early adoption -- voter rewards become claimable after round settlement, and healthy submitter rewards are snapshotted when submitter stakes return (rate halving schedule).
Consensus Subsidy	4,000,000 cREP	Pre-funded reserve for one-sided round rewards, replenished by 5% of each losing pool
Treasury	10,000,000 cREP	Governance-controlled tokens for grants, whistleblower rewards, and protocol development
Keeper Reward Pool	100,000 cREP	Flat per-operation rewards for keeper housekeeping (settlement, cancellation, cleanup), funded separately from user stakes
Category Registry	100 cREP	Initial reserve for the category proposal mechanism

4.4 HumanFaucet

Primary distribution via Self.xyz passport verification with age verification (18+). Each passport can claim once. Claim amounts decrease as more users join -- rewarding early adopters who bootstrap the platform with content.

Tier	Claimants	Claim (no referral)	Claim (with referral)	Referrer gets
0 (Genesis)	0 - 9	10,000 cREP	15,000 cREP	5,000 cREP
1 (Early Adopter)	10 - 999	1,000 cREP	1,500 cREP	500 cREP
2 (Pioneer)	1,000 - 9,999	100 cREP	150 cREP	50 cREP
3 (Explorer)	10,000 - 999,999	10 cREP	15 cREP	5 cREP
4 (Settler)	1,000,000+	1 cREP	1.5 cREP	0.5 cREP

The ~51.9M faucet pool serves up to ~41 million users without referrals (~15 million with full referral usage). Referral bonuses scale proportionally at 50% of the claim amount. The first 10 Genesis claimants receive 10,000 cREP each to bootstrap the platform from day one. As the platform grows and becomes more populated, later claimants need fewer tokens since there is already content to engage with.

4.5 Participation Pool

The participation pool solves the cold start problem. When the platform is new and vote stakes are small, round rewards alone may not be enough to attract voters and submitters. The participation pool pays proportional bonuses based on stake amount: voters claim participation rewards after round settlement regardless of vote outcome, while submitter participation rewards are only snapshotted when the submitter stake resolves on the healthy path after a settled round. The voter reward rate is snapshotted at resolution time for fairness. Early participants receive the most thanks to a halving schedule as cumulative rewards grow and the reward rate decreases.

The reward formula is:

$$\text{reward} = \text{stakeAmount} \times \text{currentRate}$$

The rate starts at 90% and halves based on cumulative cREP distributed from the pool -- making the pool's lifetime predictable regardless of individual stake sizes.

Tier	cREP Distributed	Cumulative	Rate	Stake 10	Stake 100
0	2,000,000	2,000,000	90%	9 cREP	90 cREP
1	4,000,000	6,000,000	45%	4.5 cREP	45 cREP
2	8,000,000	14,000,000	22.5%	2.25 cREP	22.5 cREP
3	16,000,000	30,000,000	11.25%	1.125 cREP	11.25 cREP

Voter participation rewards are distributed when a round settles -- deferred from vote time to prevent exploitation where attackers could vote, collect immediate participation rewards, and then have rounds cancel without risk. Submitter participation rewards are paid only when the submitter stake resolves on the healthy path after a settled round. The pool is funded with 34M cREP and governed by the same timelock as all other protocol contracts.

4.6 Keeper Network

Anyone can run a keeper. Keepers are lightweight services that monitor the blockchain for active rounds and perform reveal, settlement, reveal-failed finalization, and post-settlement cleanup work. The contract enforces a reveal grace period (1 hour) during which all past-epoch votes must be revealed before settlement is allowed, preventing selective revelation attacks. Round finalization and cleanup remain permissionless -- any account can call the relevant functions.

Keepers also perform housekeeping: cancelling expired rounds (rounds that exceed maxDuration without reaching minVoters) and marking dormant content. The drand randomness beacon is public, so anyone can run the off-chain decryption flow, but the current protocol verifies commit consistency rather than proving on-chain that the stored ciphertext was honestly decryptable. In practice the reveal path is a keeper/drand-assisted off-chain flow with a user fallback, not a fully trustless ciphertext proof system. If Curyo later wants to close that trust gap entirely, zk proofs of correct decryption are the most natural upgrade path.

To incentivize keeper operation, the protocol allocates a dedicated 100,000 cREP keeper reward pool, funded separately from user stakes. Keepers currently earn a flat 0.1 cREP for rewarded housekeeping operations such as settlement and unrevealed-vote cleanup. At this rate, the pool funds up to 1,000,000 rewarded operations.

Rewards are best-effort: if the pool is depleted, operations still succeed but no reward is paid. The keeper reward amount is governance-configurable.

4.7 Treasury

Slashed submitter stakes, the 1% treasury fee on contested settlements, and forfeited past-epoch unrevealed stakes all flow to the treasury (governance timelock). The consensus subsidy reserve is separate: it is pre-funded at launch and replenished by 5% of losing pools from two-sided rounds. Treasury tokens can only be distributed through governance proposals -- for grants, whistleblower rewards, and protocol development.

4.8 Point Distribution

When a round settles, the losing side's stakes are distributed. Winners also get their original stake back.

Recipient	Share
Revealed losing voters (rebate)	5% of raw losing stake
Winning voters (content-specific)	80% of the remaining 95%
Content submitter	10% of the remaining 95%
Consensus subsidy reserve	5% of the remaining 95%
Frontend operators	3% of the remaining 95%
Category submitter	1% of the remaining 95%
Treasury	1% of the remaining 95%

A revealed losing vote first recovers a fixed 5% rebate. The 80% voter share then goes to a content-specific pool, distributed proportionally by epoch-weighted effective stake to winning voters on that content. Tier-1 voters (who committed during epoch 1 with no information) earn full weight (100% of their stake), while Tier-2 voters (who committed after epoch-1 results were visible) earn 25% weight. This 4:1 ratio means early voters receive a larger portion of the reward pool per cREP staked. Because each content item has independent rounds that settle on their own timeline, rewards are claimable immediately after settlement -- no waiting for other content. The 5% consensus subsidy share funds one-sided-round rewards (see Consensus Subsidy Pool). The 1% treasury fee goes to the governance timelock.

4.9 Deferred Participation Rewards

Voter participation rewards are distributed at round settlement, not at vote time. This design choice eliminates a critical attack vector: if voters received an immediate participation bonus at vote time, it would reduce their at-risk capital. This could create exploitation opportunities for coordinated minorities who could stake on low-liquidity content, collect the participation reward immediately, and profit regardless of outcome.

By deferring voter rewards to settlement, the full vote stake stays at risk until the round completes. Combined with the epoch-weighted reward structure (which penalizes late entrants with 25% weight vs 100% for early voters)

and deterministic epoch-based settlement (which prevents strategic timing of entries), the deferred model ensures voter participation rewards flow only to genuine, successful curation activity while submitter bonuses unlock only after healthy settled validation.

4.10 Staking Requirements

Action	Stake	Notes
Vote on content	1-100 cREP	Per vote, per round
Submit content	10 cREP	Returned after a healthy settled round once no later round remains open, or at dormancy if no round ever settles
Register as frontend	1,000 cREP	Requires governance approval

Submitter stakes are slashed (100% to treasury) if content rating drops below 25 after a 24-hour grace period and a settled round establishes that low rating. Stakes are returned after roughly 4 days once a settled round confirms a healthy rating and no later round remains open. If no round ever settles, the stake instead resolves when the content reaches dormancy. Healthy submitter participation rewards are snapshotted at that return point and claimed later; whatever the pool can already fund is reserved immediately so later claims do not depend entirely on future authorization state.

4.11 Sybil Attack Economics

Attack Model

Consider an attacker who acquires K fraudulent verified identities at cost c per identity (passport-grade KYC). Each identity can stake up to 100 cREP per content per round, giving the attacker maximum voting power of K x 100 cREP.

Profitability Analysis

For the attack to succeed, the attacker must control the majority stake. If L_honest is the honest voters' stake on the losing side, the attacker's total winning payoff (beyond recovering stakes) is 0.76 x L_honest (80% of the post-rebate losing pool). The total cost is K x c (identity acquisition). The attack is profitable only when:

$$K < \frac{0.76 \cdot L_{\text{honest}}}{c}$$

Identity cost (c)	Honest losing stake (L)	Max profitable identities (K)
10 cREP equiv.	100 cREP	8
50 cREP equiv.	100 cREP	1
10 cREP equiv.	1,000 cREP	82

50 cREP equiv.	1,000 cREP	16

The real-world cost of a verified passport identity far exceeds any on-chain equivalent. Even at low assumed identity costs, profitability requires the attacker to control the majority -- if honest voters collectively outstake the attacker, all K identities lose their entire staked cREP. The attack is negative-sum in expectation against an active honest voter base.

Permanent Revocation Deterrent

If detected via on-chain pattern analysis (correlated wallet funding, synchronized vote timing, identical stake amounts) and a subsequent governance proposal, all K identities are permanently revoked. The attacker loses not only the current round's stake but all future voting capability across those identities. The expected cost of detection increases with K (more identities produce more on-chain correlation signals), creating a superlinear deterrent:

$$E[\text{penalty}] = P(\text{detect} \mid K) \cdot K \cdot V_{\text{future}}$$

where V_future is the discounted future value of each identity's voting participation.

5. Governance

On-chain governance for shaping the platform's future.

5.1 Overview

Curyo is designed to finalize into a fully decentralized, community-governed system. In finalized deployments, token holders govern protocol parameters on-chain and temporary setup roles are renounced so no privileged admin keys or multisigs remain. Local or pre-finalization environments may still use temporary deployer wiring during setup.

Curyo is a reputation token with no monetary value. It is not sold, has no treasury backing, and is not designed as a financial instrument. Governance power comes from earning reputation through verified participation, not from purchasing tokens.

5.2 Voting Power

Curyo includes built-in governance capabilities with snapshot-based voting. Your voting power equals your cREP balance and is activated automatically -- no delegation step required.

5.3 Proposal Lifecycle

State	Description
Pending	Created. Waiting for voting delay (~1 day / 7,200 blocks).
Active	Voting open (~1 week / 50,400 blocks). Cast: For, Against, or Abstain.
Queued	Passed. In timelock queue (2 days).
Executed	Changes are live.

5.4 Parameters

Parameter	Value
Proposal threshold	100 cREP
Voting delay	~1 day (7,200 blocks)
Voting period	~1 week (50,400 blocks)
Quorum	4% of circulating supply (min 10K cREP)
Timelock delay	2 days

Governance lock	7 days transfer-locked (after voting or proposing)

5.5 Round Voting Parameters

The following parameters control per-content round-based voting. Core round settings are adjustable via governance proposals through the setConfig() function on the RoundVotingEngine contract. The reveal grace period is updated separately through setRevealGracePeriod().

Parameter	Default	Description
epochDuration	20 minutes	Duration of each reward tier; commits in epoch 1 earn 100% weight, later epochs 25%
maxDuration	7 days	Maximum round lifetime -- below commit quorum rounds cancel; commit-quorum rounds can end as RevealFailed
minVoters	3	Minimum revealed votes required before settlement is allowed
maxVoters	1,000	Per-round cap on total commits
Rating smoothing (b_r)	50 cREP (hardcoded)	Controls rating sensitivity to individual votes
Vote stake	1-100 cREP	Stake range per vote, capped per Voter ID
Vote cooldown	24 hours	Wait time before voting on the same content again

The epoch-based mechanism ensures rounds complete within a bounded timeframe. The epochDuration defines the reward tier window (20 minutes for full weight). Settlement becomes available once minVoters is reached and past-epoch reveal constraints are satisfied. The maxDuration hard cap prevents indefinite rounds. The rating smoothing parameter b_r is hardcoded and controls how responsive the content rating is to individual revealed votes. As the platform grows, governance can adjust the configurable parameters to optimize for the observed voter population.

5.6 Treasury

The governance treasury is held by the timelock controller and starts with 10M cREP. It grows over time through three primary ongoing inflow sources:

- 1% settlement fee -- 1% of contested losing pools is sent to the treasury when rounds settle.

- Slashed submitter stakes -- when content is flagged for policy violations or receives unfavorable ratings, the submitter's 10 cREP stake is slashed to the treasury.
- Forfeited unrevealed votes -- past-epoch unrevealed stakes are swept to treasury during post-settlement cleanup.

Treasury tokens can only be distributed through governance proposals. Token holders propose allocations, the community votes, and after the timelock delay, the transaction is executed on-chain. This ensures transparent, community-controlled distribution of protocol tokens.

5.7 Collusion Prevention

The integrity of cREP's content curation depends on honest, independent voting. Groups that coordinate to artificially upvote or downvote content undermine the parimutuel system and harm fair curation.

Detection

Community members can monitor voting patterns on-chain. Suspicious activity -- such as coordinated voting from related wallets, vote timing patterns, or unusual stake distributions -- can be flagged and analyzed using on-chain data.

Enforcement via Governance Proposals

When hard evidence of collusion is found, the community can take action through governance:

- Revoke Voter IDs -- governance can permanently revoke the Voter ID NFTs of confirmed colluders, removing their ability to vote on the platform.
- Reward whistleblowers -- governance is encouraged to allocate cREP from the treasury to reward community members who provide evidence of collusion.

Deterrence

Several protocol features make collusion costly and difficult:

- Sybil resistance -- 1 person = 1 Voter ID via passport verification (Self.xyz).
- Stake caps -- maximum 100 cREP per content per round limits single-voter influence.
- Vote cooldowns -- a 24-hour cooldown on the same content prevents rapid re-voting and is enforced per effective Voter ID.
- Permanent revocation -- losing your Voter ID is irreversible and eliminates voting ability.

Formal Collusion Model

A coalition of C colluders coordinates to vote in the same direction on target content. Each colluder stakes s_c (up to 100 cREP). Their combined stake is S_C = C x s_c. Let S_H denote honest voters' stake on the opposite side. The coalition wins if S_C > S_H. Coalition profit (beyond recovering stakes) is 0.76 x S_H (80% of the post-rebate losing pool), shared among C members. Per-member profit:

$$\text{profit per member} = \frac{0.76 \cdot S_H}{C}$$

Diminishing Returns

For collusion to exceed the per-member coordination cost k (communication, trust establishment, detection risk):

$$\frac{0.76 \cdot S_H}{C} > k$$

As coalition size C grows, per-member profit shrinks linearly while coordination cost and detection risk increase. This creates a natural ceiling on profitable coalition size. Furthermore, if honest voters respond to suspected collusion by increasing their counter-stakes, S_H grows and the required coalition size increases further.

Detection Probability and Expected Penalty

On-chain signals of collusion include: identical vote timing within the same block or narrow window, correlated stake amounts, shared funding sources traceable via transaction graphs, and repeated same-direction voting on identical content across rounds. The probability of detection P(detect | C) is monotonically increasing in C. Combined with permanent Voter ID revocation, the expected penalty is:

$$E[\text{penalty}] = P(\text{detect} \mid C) \cdot C \cdot V_{\text{future}}$$

where V_future represents the net present value of each identity's future voting rewards. For sufficiently high detection probability or future voting value, the expected penalty exceeds the one-time collusion profit, making collusion a negative expected-value strategy.

The process follows cREP's standard governance flow: evidence is submitted, a governance proposal is created, the community votes, and after the timelock delay, the action is executed.

5.8 Governance Security

On-chain governance carries its own attack surface. A malicious actor who accumulates sufficient voting power could propose changes that benefit themselves at the expense of the community -- for example, altering reward splits, revoking honest Voter IDs, or draining the treasury. Curyo's governance design includes several layers of defense against such attacks.

Snapshot-Based Voting

Governance voting power is snapshot-based: it is locked at the block when a proposal is created. This prevents flash-loan attacks (borrowing tokens to vote, then returning them) and just-in-time token acquisition. An attacker must hold cREP before the proposal exists, making surprise governance attacks impractical.

Timelock Delay

All approved proposals enter a 2-day timelock queue before execution. This gives the community a window to detect malicious proposals and organize a response -- including submitting counter-proposals or alerting the broader community. The delay acts as a circuit breaker against governance capture.

Early-Stage Concentration Risk

Quorum is calculated as 4% of circulating supply -- total supply minus balances held by protocol-controlled holders excluded by the governor. In the deployment model this includes custody contracts such as HumanFaucet, ParticipationPool, RewardDistributor, RoundVotingEngine reserves, the governance timelock, and registry-held stakes. This dynamic calculation ensures governance is usable from day one: when only a small number of users have claimed tokens, the quorum scales proportionally to actual circulation rather than the full 100M supply. A minimum floor of 10,000 cREP prevents trivially small quorums in the earliest stages. As the user base grows and more tokens enter circulation, the quorum threshold increases proportionally, requiring increasingly broad consensus.

No Privileged Keys

After deployment, no admin keys, multisigs, or privileged roles exist. The timelock controller is the sole owner of all protocol contracts, and it can only execute transactions that have passed the full governance lifecycle (proposal, voting, timelock). The proposal threshold is deliberately low (100 cREP) to encourage participation -- the real protection is the combination of dynamic quorum (4% of circulating supply with a 10K cREP floor), majority vote, and timelock delay, not proposal gating. Proposal eligibility is snapshot-based, so the same voting power can back multiple concurrent proposals, and the 7-day governance lock does not add marginal collateral per live proposal.

6. Curyo & AI

How stake-weighted curation addresses the AI content crisis and produces public quality infrastructure.

6.1 The Model Collapse Problem

Research by Shumailov et al. (Nature, 2024) demonstrates that AI models trained recursively on AI-generated content undergo 'model collapse' -- a progressive loss of distributional fidelity where each successive generation of models loses the tails of the original data distribution. As AI-generated content proliferates across the web, the training data available to future models becomes increasingly contaminated with synthetic output, accelerating this degradation cycle.

The implication is that verified human quality signals become critical infrastructure for maintaining the fidelity of AI systems. Without reliable mechanisms to distinguish high-quality content from low-quality or AI-generated filler, training pipelines face an increasingly noisy signal-to-noise ratio. Curyo addresses this by producing stake-weighted, Sybil-resistant quality ratings anchored to economic commitment from verified human identities.

6.2 Stake-Weighted Curation

The concept of 'staked media' (a16z, Big Ideas 2026, https://a16z.com/newsletter/big-ideas-2026-part-3/#the-rise-of-staked-media) -- where content quality is assessed through economic commitment rather than algorithmic engagement -- provides a manipulation-resistant alternative to traditional curation mechanisms. Curyo implements this approach through its parimutuel voting system: voters stake cREP tokens on their quality predictions, and the tlock commit-reveal scheme combined with epoch-weighted rewards ensures economic independence by hiding votes during epoch 1 and penalizing late herders with 25% reward weight.

This design produces quality signals with several properties that distinguish them from engagement-based metrics:

- Economic commitment -- Each rating is backed by a token stake, making systematic manipulation expensive relative to the signal produced.
- Economic independence -- tlock encryption hides votes during epoch 1, eliminating herd signals. Epoch-weighted rewards (4:1 ratio) further penalize late followers, incentivizing genuine early assessment over copying.
- Sybil resistance -- Passport-verified Voter IDs limit each human to one identity with a capped stake per content, preventing bot farms from flooding the signal.
- Verifiability -- All votes, stakes, and outcomes are recorded on-chain with cryptographic integrity, enabling third-party audit and reproducibility.

6.3 On-Chain Ratings as Public Infrastructure

A foundational design decision in Curyo is the use of a public blockchain as the settlement layer. This ensures that all quality ratings -- including individual vote directions, stake amounts, round outcomes, and resulting content scores -- are inherently public, permissionless, and exportable. No API key, rate limit, or terms-of-service restriction mediates access to the data.

This positions Curyo's output as public goods infrastructure rather than a proprietary dataset:

- AI training pipelines can incorporate on-chain quality scores to filter or weight training corpora, mitigating model collapse by prioritizing human-verified content.
- Search engines and recommendation systems can consume on-chain ratings as an independent quality signal, reducing dependence on engagement-based proxies.
- Researchers retain full transparency into voting dynamics, curation patterns, and content quality trends without data access barriers.
- Third-party platforms can build on the quality layer without permission, payment, or partnership agreements.

Unlike centralized rating platforms where data is siloed behind proprietary APIs, blockchain-native ratings function as a commons. This aligns with the thesis that the AI-dominated web requires open, verifiable quality infrastructure rather than additional walled gardens.

6.4 AI-Assisted Voting

Curyo incorporates AI as a first-class participant through automated voting bots that use pluggable rating strategies. Each strategy queries an external API to obtain a normalized quality score for submitted content. The bot votes UP or DOWN based on whether the score meets a configurable threshold.

Submission bots can also publish richer metadata than a single free-form caption. They submit a short title, a longer description, tags, and a category alongside the canonical URL, which makes downstream discovery interfaces easier to scan while keeping the same shared on-chain event history for every frontend.

Bots use the same transferAndCall-based vote commit flow as human voters and participate under the same tlock privacy constraints -- their vote direction is hidden until the epoch ends, just like human votes. Bots stake the minimum amount of cREP per vote, ensuring their influence remains small relative to human voters who may stake significantly more. Voting in epoch 1 (before any results are visible) gives bots the same 100% reward weight as early human voters, rewarding accurate strategies. The parimutuel mechanism provides natural selection pressure: strategies that produce inaccurate ratings lose their stakes, while accurate strategies accumulate reputation.

Human Oversight

The system is designed so that human voters retain decisive influence. Bots staking the minimum are outweighed by any human voter staking more. In contentious rounds, the aggregate human stake dominates bot contributions. This creates a hybrid model: AI provides baseline signals and seeding, while humans provide authoritative quality judgments.

Cold-Start Mitigation

AI-assisted voting directly addresses the cold-start problem inherent in new content platforms. When a content item is submitted, automated strategies can produce initial quality signals within seconds, seeding the voting market before human participants engage. This creates immediate activity and provides a focal point for human voters to agree or disagree with, accelerating convergence toward accurate ratings.

6.5 Future Directions

Curyo's architecture enables several extensions at the intersection of AI and decentralized curation:

- Cross-platform quality oracle -- On-chain content ratings can serve as an oracle for other protocols and platforms, creating a shared quality layer across the decentralized web.
- Expertise-weighted reputation -- Domain-specific reputation multipliers could allow voters with demonstrated accuracy in specific categories to earn additional influence, improving signal quality in specialized domains.
- Content provenance integration -- Combining Curyo ratings with content provenance standards (C2PA) would create a two-layered trust system: provenance verifies origin, stake-weighted curation verifies quality.
- Advanced AI strategies -- The pluggable strategy interface supports increasingly sophisticated approaches, from API-based lookups to LLM-driven content analysis. The parimutuel mechanism ensures that only strategies producing accurate ratings survive long-term.

7. Known Limitations

Transparency about design trade-offs, residual risks, and areas for improvement.

7.1 drand Network Dependency

tlock encryption relies on the drand randomness beacon network to produce the decryption key after each epoch ends. If the drand network experiences downtime, newly committed votes cannot be revealed until drand resumes. In practice, drand operates across a globally distributed set of nodes (the League of Entropy) and has maintained high availability since 2019. Additionally, any party who already knows the plaintext `(isUp, salt)` for a commit can manually call revealVoteByCommitKey() once the epoch ends; connected users can do this from the fallback UI. Rounds are not cancelled due to drand downtime -- they simply wait for reveals and settle once conditions are met.

7.2 tlock Reveal Burden

Although the keeper reveals votes automatically in the background, the protocol-level reveal function can also be called directly by a voter who knows the plaintext `(isUp, salt)` for their commit. The current default UX remains keeper-driven automatic reveal, but the production UI now exposes a small fallback page for connected users. That page decrypts the on-chain ciphertext locally after epoch end rather than persisting reveal secrets in browser storage by default, because long-lived localStorage copies would increase the blast radius of any frontend XSS bug.

7.3 Consensus Subsidy Pool

The parimutuel reward structure distributes the losing pool to winners. When all voters vote in the same direction (one-sided round), the losing pool is zero and no standard parimutuel rewards are distributed. Without mitigation, this creates a perverse incentive: no reason to vote on obviously good or bad content, and coordinated groups could benefit from manufacturing dissent by having one member vote against the majority to create a losing pool.

The consensus subsidy pool solves this. It is pre-funded with 4,000,000 cREP from the treasury allocation and continuously replenished by 5% of every losing pool from two-sided rounds. When a one-sided round settles (all votes in the same direction), the contract distributes a subsidy from this reserve equal to 5% of the round's total stake, capped at 50 cREP per round and by the reserve balance.

The subsidy formula is:

$$\text{subsidy} = \min\left(\text{reserveBalance},\ \text{totalStake} \times 0.05\right)$$

This subsidy is split between voters (~89%) and the content submitter (~11%), using the same 82:10 ratio as normal round rewards, and distributed proportionally by epoch-weighted effective stake within each group. Since all voters are on the winning side, every voter receives a share. The mechanism is self-sustaining: contentious rounds -- where parimutuel rewards function normally -- generate surplus that funds consensus rounds. Every two-sided round with L cREP in its losing pool contributes 0.05L to the reserve, which can fund approximately one one-sided round of equivalent total stake. The 4M initial pre-fund provides runway during early adoption when two-sided rounds may be infrequent.

Consensus subsidy rewards are intentionally lower than contentious-round rewards (approximately 10:1 ratio), preserving the incentive to vote on genuinely contentious content while making consensus curation non-zero.

7.4 Configuration Change Timing

Governance can change round parameters (epochDuration, maxDuration, minVoters) at any time through the standard proposal process. Changes apply to new rounds only: each round snapshots configuration at creation time, so in-progress rounds keep the rules they started with.

7.5 Settlement External Dependencies

Round settlement interacts with external contracts (ParticipationPool, FrontendRegistry, CategoryRegistry) using fail-soft wrappers for non-critical side effects. If one of these external calls reverts, settlement can continue while skipping that side effect, preventing total settlement blockage at the cost of temporarily deferred accounting or payouts for that component.

7.6 Identity Verification Scope

Passport-based identity verification via Self.xyz provides strong Sybil resistance but excludes approximately 1.1 billion people globally who lack passports. The system has no appeal mechanism for false rejections, and recovery from a compromised or offline Self.xyz service is not documented. These are inherent trade-offs of passport-gated identity systems.