Reppo: Tokenized Infrastructure for Reinforcement Learning via Prediction Markets

Kiran Karra, Jordan Grollman, Raghav Rmadya

Reppo Foundation

Abstract

AI systems are limited not by compute, but by access to continuously updated, high-quality training data. Reinforcement Learning (RL) has long been the dominant approach to training AI models and systems based on human feedback and experience rather than historical data. In the RL domain, a model takes an action, usually in a game or simulator, and then receives feedback on whether those actions are productive in helping it achieve a goal. Through trial and error over the course of many actions, the AI learns the best ways to achieve the goal. Today's RLHF pipelines rely on static datasets and loosely aligned human feedback. They lack mechanisms to price disagreement or ensure accountability in data curation. Reppo introduces a new primitive: data quality as economically weighted consensus. By combining reinforcement learning with prediction markets, Reppo turns data labeling into a capital-backed process where disagreement is priced and becomes part of the training signal.

> Data quality = economically weighted consensus through prediction markets"

In Reppo, **Datanets** function as authorized data companies or entities on the chain. Each Datanet is a two-sided prediction market where Publishers, i.e. data contributors, publish raw data on one, and subject matter experts lock $REPPO to provide their opinions and feedback on the published data through a prediction market mechanics. The data produced are curated and continuously updated, ready to be consumed downstream by humans, agents, models, and robotics systems. Reppo does not claim consensus equals objective truth. Instead, it provides economically weighted belief. Over multiple epochs, with diverse participants and capital at risk, this belief converges toward high-confidence signals. This paper outlines the Datanet architecture, the vote-escrowed governance model (veREPPO), the net-vote and linear-decay voting mechanics, the emission and revenue flows that sustain each Datanet economy, and the resulting flywheel that ties token value to data quality at the protocol level.

1 Background and Motivations

Reinforcement learning has long been the dominant approach to training AI models and systems based on human feedback rather than historical data and it has driven some of the most consequential results in AI—from AlphaGo to large-scale language model alignment via RLHF (Reinforcement Learning from Human Feedback). As we enter an era defined by real time experience - AI models, agents and robots learning from grounded real-world interactions rather than simulated reward signals, traditional RL pipelines need an upgrade. Current RLHF approaches lack aligned incentives between the various actors involved in the process, often leading to costly mistakes that, in some cases, can lead to AI that is biased, misaligned, wrong, or shortsighted. To build true superintelligence, AI must learn to achieve its goals from scratch, often working from first principles. A growing body of work, led by researchers such as David Silver and Rich Sutton, argues that the next frontier is *experiential* RL: agents that learn predominantly from grounded interaction with real environments, supplemented by high-signal human feedback as a bootstrap [1].

Two problems stand in the way:

1. **Data staleness.** Most training datasets are static snapshots. Real–world tasks, such as robotic manipulation, autonomous navigation, use of agents, require data that updates continuously as environments and human preferences shift.

2. **Quality assurance without centralized oracles.** Platforms like Scale AI rely on low-wage annotation pipelines with limited economic accountability. There is no on-chain mechanism for crowd-sourced annotation & labeling, where data curators wager capital at risk behind their assessments.

Decentralized AI projects have made progress on supply-side incentives for compute and storage, but remain stuck in partnership-driven distribution with limited organic demand. Meanwhile, AI agents and physical AI systems need permissionless, real-time access to curated data regardless of supplier or ecosystem lock-in.

Reppo addresses these gaps by treating data curation as a staked prediction market and data consumption as a subscription service.

2 Introduction

The Reppo protocol rests on two core mechanisms:

- **Datanets are on-chain data businesses;** Each Datanet is a self-contained data market –represented by an on-chain NFT –where a Datanet owner defines a data collection goal, sets publishing and access fees, and seeds an emission pool to incentivize contributors. Publishers submit raw data (robotics trajectories, video, images, text, code); voters lock REPPO to evaluate it. The result is a continuously updated, curated data set that the Datanet owner monetizes through the Reppo Data Exchange.

- **Vote-escrowed tokenomics with stake-assured feedback.** Reppo extends the vote-escrowed tokenomics (veTokenomics) model so that REPPO holders who lock tokens receive voting power used to assess data quality within datanets. Unlike traditional veTokenomics, where votes direct emissions to projects, here votes serve a dual purpose: they direct economic rewards *and* they produce a quality signal—a synthetic prediction market where consensus is the oracle.

This paper focuses on the Datanet architecture, the voting and emission mechanics, and the economic model that supports the protocol. A separate technical overview covering the Data Exchange and Infrastructure Exchange will be released later.

3 Core Concepts and Stakeholders

3.1 Datanets

A **Datanet** is an on-chain data business. Each Datanet target a specific data domain – tokenomics research, robotic grasping trajectories, street-level imagery, code quality assessments – and provides the economic scaffolding for crowdsourced data collection and curation.

datanets are represented as NFTs. The owner of the NFT controls the Datanet's parameters: publishing fees, access gating, emission pool funding, and the split of rewards between publishers and voters. Because the Datanet is an NFT, ownership is transferable—a Datanet generating consistent revenue can be sold, separating the roles of Datanet creator and Datanet operator.

Each Datanet defines:

- A **Datanet goal**: a description of the target data set and its intended use.
- **Publisher instructions**: criteria for what raw data should be submitted.
- **Voter instructions**: evaluation criteria for assessing published content.
- **Fee structure**: the publishing fee charged to data contributors and any access fee for gated entry.
- **Emission configuration**: token and amount seeded into the reward pool and split ratio between publishers and voters.

3.2 Pods

A **Pod** is a unit of raw content published in a Datanet. When a publisher submits data—a paper, an image set, a trajectory recording, a code snippet—it becomes a Pod within that Datanet. Pods are the objects that voters evaluate. Each Pod accumulates positive and negative votes over the course of an epoch. Pods with positive net votes earn emission rewards for their publisher and supporting voters; Pods with negative net votes are filtered from the curated dataset and receive no emissions.

3.3 Stakeholders

There are four participant roles in the Reppo ecosystem:

- **Datanet Owners:** Create and operate datanets. They pay a creation fee to the network, seed the emission pool with tokens, set publishing and access fees, and earn revenue from those fees. Their goal is to build a high-quality, continuously updating dataset and monetize it through the Reppo Data Exchange or direct access subscriptions.
- **Publishers:** Contribute raw data (Pods) to datanets. Publishers pay a publishing fee set by the Datanet owner, and in return, earn a share of the Datanet's emission pool proportional to the positive net votes their Pods receive.
- **Voters (veREPPO Holders):** Lock $REPPO tokens to receive veREPPO voting power, then allocate that power to assess Pods within datanets –voting positively or negatively. Positive voters earn emissions when the Pods they support achieve positive consensus. Negative voters serve a filtering role, preventing low-quality content from accruing rewards. This is the mechanism we call *stake-assured human feedback*: quality judgments backed by locked capital.
- **Data Consumers:** Subscribe to curated datasets through the Reppo Data Exchange (repo.exchange). Because Datanet data are updated every epoch (48 hours), consumers receive a continuous data pipeline rather than a static snapshot, the core value proposition for downstream RL training, fine-tuning, and agent development.

3.4 Reppo Token Utility

The $REPPO token serves several functions within the ecosystem:

- **Governance (veREPPO):** $REPPO tokens locked by holders become veREPPO, which provides voting power to assess data quality within datanets and to participate in protocol-level governance decisions.
- **Datanet Creation:** Creating a Datanet requires paying a fee denominated in $REPPO, a portion of which is burned and a portion locked.
- **Publishing and Access:** Datanet owners may denominate publishing and access fees in $REPPO or other tokens. When non-REPPO tokens are used

for emission seeding, the protocol levies a tax, creating buy pressure on $REPPO.

- **Data Exchange:** $REPPO serves as a payment and reward token on the Data Exchange (repo.exchange) and Infrastructure Exchange (infra.exchange), with stablecoin settlement also supported.

4 Technical Architecture

The Reppo protocol consists of several on-chain and off-chain components that coordinate Datanet creation, data publishing, voting, emission distribution, and data exchange.

4.1 Datanet Manager

The Datanet Manager governs the lifecycle of datanets on the network.

4.1.1 Datanet Creation (V2)

In V2, Datanet creation is fully self-served. A prospective Datanet owner pays a creation fee—currently targeted at 10,000–20,000 $REPPO, set by governance via Snapshot. This fee is divided as follows:

- **Burn (∼10%):** Permanently removed from circulation.
- **Incentivization Pool (∼40%):** Allocated to a protocol-level pool that rewards veREPPO holders who back datanets (see Section 6.4).
- **Lock (remainder):** Returned to the Datanet owner when the Datanet is wound down. Winding down requires burning the Datanet NFT.

Upon payment, the protocol mints a Datanet NFT to the creator. This NFT represents ownership and control of the Datanet. The owner can transfer or sell this NFT at any time, enabling Datanet businesses to change hands.

4.1.2 Emission Seeding

Once a Datanet is live, the owner seeds an emission pool to incentivize publishers and voters. This is the operating cost of running a Datanet. The owner can seed emissions in:

- **$REPPO:** No tax applied.
- **Any other token:** A governance set tax is deducted and sent to the protocol treasury for operations and buybacks. The remainder enters the Datanet's emission pool.

This design choice means Reppo does not rely on inflationary protocol-level emissions. Datanet owners fund their own incentive pools, and the protocol captures value through the creation fee burn and the non-REPPO emission tax.

4.1.3 Datanet Revenue

Datanet owners generate revenue from two on-chain sources and one off-chain source:

1. **Publishing fees:** Charged to every publisher who submits a Pod. The Datanet owner sets the amount and denomination. The protocol takes a % cut.
2. **Access fees:** Charged to voters or viewers who enter a gated Datanet. The Datanet owner sets the amount.
3. **Data Exchange revenue:** Datanet owners list their curated datasets on repo.exchange, where data consumers subscribe. Since datasets update every epoch, the subscription model produces recurring revenue.

The Datanet owner's profit is the difference between revenue (publishing fees + access fees + exchange subscriptions) and cost (emission seeding). A well-run Datanet can be cash-flow positive from on-chain fees alone, with exchange revenue as upside.

4.2 Voting Mechanics (V2)

V2 introduces three changes to the voting system: net voting, linear decay, and the removal of commit-reveal.

4.2.1 Net-Vote Model

Voters can now cast both **positive** and **negative** votes in pods. This transforms the system from a simple ranking mechanism into a prediction market where disagreement carries a signal.

For each Pod j in a given epoch, define:

$$\mathrm{NetVote}_j = \sum_{u \in U_j^+} v_{u,j} - \sum_{u \in U_j^-} v_{u,j} \tag{1}$$

where U_j^+ and U_j^- are the sets of positive and negative voters for Pod j, and $v_{u,j}$ is the effective voting power allocated by user u.

The payout logic depends on the sign of $\mathrm{NetVote}_j$:

- **Positive net vote ($\mathbf{NetVote}_j > 0$):** Emissions allocated to Pod j are split between the publisher and positive voters according to the configured split ratio of Datanet. Negative voters receive nothing.

- **Negative net vote ($\mathbf{NetVote}_j < 0$):** The publisher is zeroed out—they receive no emissions for that Pod. The emissions that would have been allocated to Pod j remain in the pool and are redistributed among Pods with positive net votes.
- **Zero net vote ($\mathbf{NetVote}_j = 0$):** No emissions are distributed for that Pod.

This mechanism captures disagreement as a data signal. For downstream consumers of AI training data, knowing which content generated a strong negative consensus is as valuable as knowing which content was endorsed.

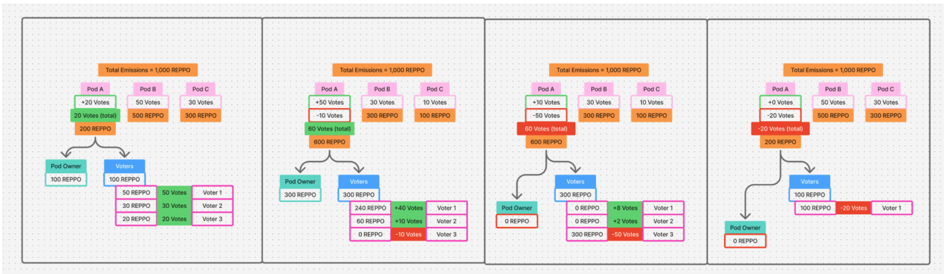

Figure 1: Net-vote payout scenarios. Four cases illustrating how positive and negative votes interact to determine emission distribution between publishers and voters.

4.2.2 Linear Decaying Voting Power

To reward early conviction and prevent last-minute manipulation, voting power decays linearly over the epoch. A voter who commits at the start of an epoch uses their full veREPPO power; a voter who commits at the midpoint uses half.

Formally, if an epoch spans time $[0, T]$ and a voter casts their vote at time t:

$$v_{u,j}(t) = V_u \cdot w_{u,j} \cdot \left(1 - \frac{t}{T}\right) \tag{2}$$

where V_u is the total veREPPO power of the voter and $w_{u,j}$ is the fraction allocated to Pod j.

This decay curve incentivizes voters to evaluate content early in the epoch when less information is available, rewarding genuine assessment over strategic pile-on.

4.2.3 Removal of Commit-Reveal

V1 used a commit-reveal scheme: voters committed hidden votes in the first phase of an epoch, then revealed them in the second phase. Only revealed votes counted toward settlement.

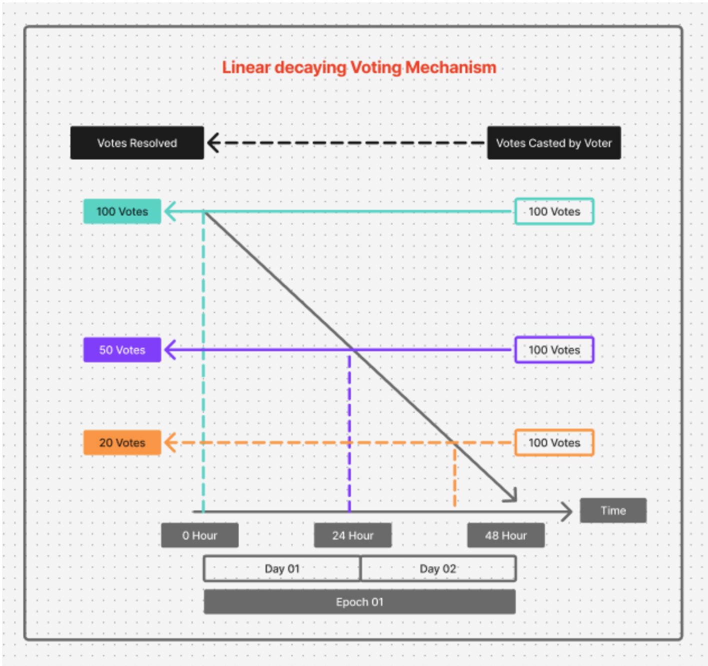

Figure 2: Linear decaying voting mechanism. A voter casting 100 votes at hour 0 retains full power (100 resolved); the same 100 votes cast at hour 24 resolve to 50; at hour 48 they approach zero.

V2 removes this mechanism entirely. Votes are visible in real time. This change serves two purposes:

1. **Transparency:** Publishers and voters can observe which Pods are leading or trailing, increasing engagement and veREPPO volume—a key network health metric.
2. **Reduced friction:** The reveal step introduced a second required transaction that many users failed to complete, reducing the effective voter pool.

The linear decay mechanism compensates for the loss of hidden voting by making early votes more valuable, preserving the incentive for independent judgment.

4.2.4 Synthetic Oracle and Market Resolution

Traditional prediction markets resolve against an external oracle—an event occurs or it does not. Reppo's prediction market resolves against internal consensus: the net vote tally at the end of the 48-hour epoch *is* the oracle. We term this a **synthetic oracle**.

Resolution occurs deterministically at the close of every epoch. At the epoch boundary:

1. All votes are finalized and the linear decay weights are locked.
2. For each Pod, NetVote_j is computed from the decay-weighted sum of positive minus negative votes.
3. The sign of NetVote_j determines the market outcome: positive net vote means the market resolves in favor of the publisher (Pod accepted into the curated dataset); negative net vote means the market resolves against (Pod excluded, publisher zeroed out).
4. Emissions are distributed according to the resolved outcomes (Section 6).
5. The curated dataset state is updated—accepted Pods are included, rejected Pods are filtered out—and the new dataset version is pushed to subscribers on repo.exchange.

No external arbiter, oracle feed, or manual adjudication is required. The resolution is fully endogenous: the crowd of economically committed voters *is* the oracle. Over multiple epochs (typically six to seven), the repeated consensus converges toward a high-confidence quality assessment for the Datanet's dataset. This is the mechanism that transforms raw crowdsourced data into curated training data suitable for RL pipelines.

4.3 Epoch Structure

All Datanet activity is organized into 48-hour epochs. The lifecycle of each epoch proceeds as follows:

1. **Publishing phase:** Publishers submit Pods (raw data) and pay publishing fees. New Pods become eligible for voting immediately upon submission.
2. **Voting phase (concurrent):** Voters allocate veREPPO power to Pods (positive or negative). Votes are visible in real time. Linear decay applies—earlier votes carry more weight.
3. **Settlement:** At epoch close, the synthetic oracle resolves all markets. Net votes are tallied, emissions are distributed, and Pods are classified as accepted or rejected.
4. **Dataset update:** The curated dataset state is updated and the new version is available to subscribers on repo.exchange within minutes of settlement.

4.4 Governance Framework

Governance operates at two levels:

Datanet-Level Governance. Datanet owners control parameters specific to their Datanet: publishing fees, access fees, emission pool funding, the publisher/voter split ratio, and content guidelines. This autonomy allows each Datanet to optimize for its specific data domain and business model.

Protocol-Level Governance. The Reppo DAO, operating through Snapshot, governs system-wide parameters:

- Datanet creation fee (amount and burn/lock/incentivization split)
- Non-REPPO emission tax rate
- Protocol cut on publishing fees
- Epoch duration
- Incentivization pool distribution parameters

5 Token Mechanics

5.1 Vote-Escrowed Token System (veREPPO)

veREPPO is the governance and voting primitive of the Reppo protocol. When users lock their $REPPO tokens, they receive veREPPO in return, with voting power proportional to both the quantity locked and the duration of the lock.

The relationship between lock duration and voting power follows a non-linear function:

$$V_u = X_u \cdot \left(\frac{L_u}{L_{max}}\right)^{\gamma} \tag{3}$$

Here, X_u is the amount of $REPPO locked, L_u is the chosen lock duration in epochs, L_{max} is the maximum allowed lock duration, and $\gamma > 1$ is a parameter controlling non-linearity. Longer lock durations yield disproportionately more voting power, incentivizing sustained commitment to the network.

Voting power can be reallocated across datanets and Pods each epoch, providing capital allocation flexibility despite the long-term lock commitment.

5.2 Voting Power Allocation

veREPPO holders deploy their voting power within datanets. For each epoch, a voter decides:

1. Which Datanet(s) to participate in.
2. Which Pod(s) within those datanets to vote on.

3. Whether to vote positive or negative on each Pod.
4. How much voting power to allocate to each vote.

The effective voting power applied to each vote is subject to the linear decay described in Section 4.2.

5.3 Locking REPPO vs. Staking

A terminological note: throughout this paper, we use "locking" rather than "staking" to describe the act of depositing $REPPO into the veREPPO contract. Locking denotes a time-bounded commitment that grants voting rights. The locked tokens are not delegated to a validator, not placed at risk of slashing, and not pooled for yield generation. The distinction matters both technically and from a regulatory perspective. Section 6.4 introduces a separate Datanet assessment mechanism—predicting Datanet performance—which is distinct from the veREPPO lock.

6 Economic Model

The Reppo economic model operates at two levels: the Datanet economy (micro) and the protocol economy (macro).

6.1 Datanet Economics

Each Datanet is an independent data business with its own P&L:

Revenue	**Cost**
Publishing fees (per Pod submitted)	Emission pool seeding
Access fees (gated datanets)	Datanet creation fee (one-time)
Data Exchange subscriptions	Protocol tax on non-REPPO seeding

Table 1: Datanet Owner P&L Structure

The Datanet owner's objective is straightforward: attract enough publishers and voters to produce a high-quality dataset, then monetize that dataset through the Data Exchange. Publishing and access fees provide near-term revenue; exchange subscriptions provide recurring revenue as the dataset matures.

6.1.1 Emission Distribution Within a Datanet

Each epoch, the Datanet's emission pool distributes rewards to publishers and voters of Pods with non-zero net votes. The distribution follows these rules:

1. Compute the share of total positive net votes attributable to each Pod:

$$s_j = \frac{\max(\text{NetVote}_j, 0)}{\sum_k \max(\text{NetVote}_k, 0)} \tag{4}$$

2. For Pods with $\text{NetVote}_j > 0$, allocate $s_j \cdot E_{pool}$ tokens, where E_{pool} is the epoch's emission pool. Split this allocation between the publisher and the positive voters according to the configured Datanet ratio ρ (default 50/50):

$$\text{PublisherReward}_j = \rho \cdot s_j \cdot E_{pool} \tag{5}$$

$$\text{VoterReward}_j = (1 - \rho) \cdot s_j \cdot E_{pool} \tag{6}$$

   Voter rewards are distributed proportionally to each voter's effective veREPPO power (after decay) assigned to Pod j.

3. For Pods with $\text{NetVote}_j < 0$, the publisher receives nothing. The emissions that would have been allocated to that Pod remain in the pool and are redistributed across positive-net-vote Pods. Negative voters do not receive emissions directly; their incentive is to prevent low-quality data from entering the curated dataset.

6.1.2 Publisher/Voter Split Configuration

The Datanet owner configures the split ratio ρ according to their strategy:

- A Datanet that needs to attract raw data (e.g., crowdsourcing weather observations) may set ρ high, directing more emissions to publishers.
- A Datanet where the owner publishes their own data and only needs curation (e.g., a project seeking community feedback on their outputs) can be set $\rho = 0$, directing all emissions to voters.

6.2 Third-party incentives Mechanism

Reppo implements a third-party incentive mechanism, similar to the bribing mechanism in concept to Aerodrome. External parties who are supportive of a Datanet's potential can contribute tokens to its emission pool. In exchange, they receive a share of the fees the Datanet owner earns from publishing and access. This allows Datanet owners to reduce their upfront seeding costs and gives external capital a way to participate in promising data markets without operating a Datanet themselves.

The mechanism also creates buy pressure on $REPPO: contributors who seed in non-REPPO tokens pay the emissions tax, while seeding in $REPPO is tax-free.

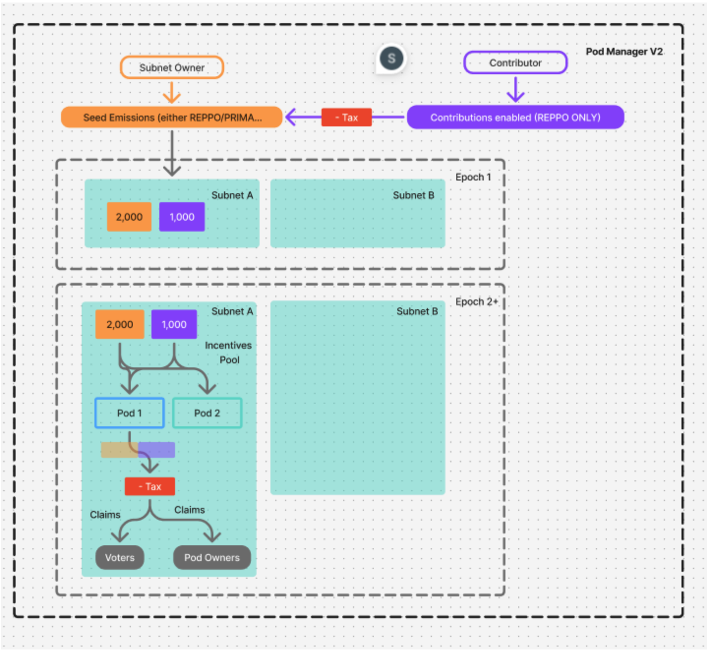

Figure 3: Emissions distribution flow. The Datanet owner seeds emissions (in REPPO or other tokens, with tax applied to non-REPPO). Emissions flow into the Datanet's incentives pool, then distribute to Pods based on net votes, with claims split between voters and publishers.

6.3 Protocol Economics

At the protocol level, Reppo captures value through several channels:

1. **Datanet creation fee burn:** ∼% of every Datanet creation fee is permanently burned.

2. **Non-REPPO emission tax:** % of all non-REPPO tokens used to seed emissions flows to the protocol treasury.

3. **Cut in publication fee:** % of all publishing fees in all datanets flows into the protocol.

4. **Data Exchange fees:** Transaction fees on repo.exchange flow back to $REPPO holders.

In particular, Reppo does not rely on inflationary emissions at the protocol-level. There is no emission mining in the style of earlier decentralized AI protocols. Datanet owners fund their own incentive pools, making $REPPO structurally deflationary – tokens are burned on Datanet creation and bought on the open market by Datanet owners who need them.

6.4 Datanet Performance Assessment

To provide an additional quality signal and create cross-Datanet competition, the protocol allows veREPPO holders to back specific datanets, predicting which will generate the most activity (measured by veREPPO voting volume, publishing volume, and exchange revenue) in the next epoch.

Participants who correctly predict high-performing datanets earn rewards from the protocol-level incentivization pool (funded by the 40% allocation from Datanet creation fees). This mechanism creates a monitoring layer: participants have an economic incentive to evaluate Datanet quality and direct attention toward productive datanets.

6.5 Economic Value of Feedback (EVOF)

A persistent question in crowdsourced data systems is: *how does a Datanet owner or dataset consumer know when the curated data is "good enough"?* The challenge applies equally to Datanet owners deciding when to stop seeding emissions and to downstream buyers evaluating whether a dataset is reliable for production use.

We propose the **Economic Value of Feedback (EVOF)** – a proprietary confidence metric that measures the average economically credible conviction behind each unit of feedback in a Datanet. EVOF answers: "How much real money-backed conviction exists behind this dataset?"

6.5.1 Definition

The EVOF is derived from the vote power-weighted average of feedback activity. Define the weighted average conviction per vote as follows:

$$W = \log(N) \cdot \frac{\sum_{i=1}^{N} \sqrt{VP_i} \cdot \text{votes}_i}{V} \tag{7}$$

where VP_i is the voting power of participant i (derived from stake amount and lock duration via the veREPPO formula), votes_i is the number of votes cast by participant i, V is the total number of votes in the epoch, and N is the number of unique voters. Then EVOF is:

$$\text{EVOF} = W^2 \tag{8}$$

The formulation captures three properties simultaneously:

1. **Economic conviction:** The $\sqrt{VP_i}$ term incorporates time-weighted economic commitment (stake size $\times$ lock duration) into every vote. Higher voting power increases influence but with diminishing returns via the square root, preventing disproportionate dominance by large stakeholders.
2. **Participation breadth:** The $\log(N)$ factor provides a soft reward for voter diversity. A Datanet with 4,000 unique voters scores higher than one with 3,000 voters at equivalent total volume, reflecting the intuition that broader participation produces more robust consensus.
3. **Interpretability:** Squaring W restores the result to a voting-power scale, making EVOF directly comparable across datanets and across time.

6.5.2 Benchmarking and Confidence Signals

EVOF is computed at two levels:

- **Datanet-level EVOF:** Computed per Datanet per epoch. This gives Datanet owners a live signal of how much economic conviction underlies their curated dataset.
- **Network-level EVOF:** The aggregate EVOF across all datanets, serving as a benchmark. A Datanet whose EVOF consistently exceeds the network average has stronger economic backing behind its quality assessments.

By tracking EVOF as a time series, Datanet owners can identify the point at which additional emission seeding yields diminishing returns in feedback quality—the inflection point where the dataset is "good enough" to reduce or stop emissions. Dataset consumers can compare a Datanet's EVOF trajectory against the network benchmark to assess whether the curation process has sufficient economic depth to trust.

6.5.3 Pairing EVOF with Datanet Performance

EVOF serves as a *preliminary* quality proxy. For a more complete picture, it can be paired with the Datanet performance metrics used in the assessment mechanism (Section 6.4):

- Total publishing volume and access fee revenue
- Total voter volume (veREPPO deployed)
- On-chain revenue generated from data monetization on the Data Exchange

A Datanet with high EVOF *and* growing exchange revenue provides the strongest confidence signal: the data has both deep economic backing from curators and demonstrated demand from consumers. The live EVOF confidence graph, displayed at the Datanet level, gives owners a real-time benchmark for how their Datanet performs relative to comparable datanets across the network.

7 Mathematical Formulation

This section consolidates the mathematical representation of the Reppo V2 model.

7.1 Notation and Parameters

- E – An epoch (48-hour period)
- T – Duration of an epoch in time units
- E_{pool} – Total emission pool for a Datanet in epoch E
- L_{max} – Maximum lock duration for veREPPO
- γ – Non-linearity parameter for veREPPO power ($\gamma > 1$)
- ρ – Publisher/voter split ratio (set by Datanet owner, $0 \le \rho \le 1$)
- fee_{create} – Datanet creation fee in $REPPO
- β_{burn} – Fraction of creation fee burned ($\sim$0.10)
- β_{pool} – Fraction of creation fee to incentivization pool ($\sim$0.40)
- β_{lock} – Fraction of creation fee locked ($= 1 - \beta_{burn} - \beta_{pool}$)
- τ_{tax} – Tax rate on non-REPPO emission seeding ($\sim$

7.2 veREPPO Voting Power

For user u who locks X_u tokens for L_u epochs:

$$V_u = X_u \cdot \left(\frac{L_u}{L_{max}}\right)^{\gamma} \tag{9}$$

7.3 Linear Decay

For a vote cast at time t within an epoch spanning $[0, T]$:

$$v_{u,j}(t) = V_u \cdot w_{u,j} \cdot \left(1 - \frac{t}{T}\right) \tag{10}$$

where $w_{u,j}$ is the fraction of voting power user u allocates to Pod j.

7.4 Net Vote Calculation

For Pod j in epoch E:

$$\text{NetVote}_{j,E} = \sum_{u \in U_j^+} v_{u,j}(t_u) \; - \; \sum_{u \in U_j^-} v_{u,j}(t_u) \tag{11}$$

where t_u is the time at which user u cast their vote.

7.5 Emission Allocation

Define the positive-net-vote share for Pod j:

$$s_{j,E} = \frac{\max(\text{NetVote}_{j,E},\ 0)}{\sum_{k \in \mathcal{P}} \max(\text{NetVote}_{k,E},\ 0)} \tag{12}$$

where $\mathcal{P}$ is the set of all Pods in the Datanet during epoch E.
For Pods with $\text{NetVote}_{j,E} > 0$:

$$\text{PublisherReward}_{j,E} = \rho \cdot s_{j,E} \cdot E_{pool} \tag{13}$$

$$\text{PositiveVoterReward}_{j,E} = (1 - \rho) \cdot s_{j,E} \cdot E_{pool} \tag{14}$$

Individual voter u's share of the positive voter reward:

$$r_{u,j,E} = \text{PositiveVoterReward}_{j,E} \cdot \frac{v_{u,j}(t_u)}{\sum_{u' \in U_j^+} v_{u',j}(t_{u'})} \tag{15}$$

For Pods with $\text{NetVote}_{j,E} < 0$, the publisher receives zero. Because $s_{j,E} = 0$ for these Pods, their would-be allocation is implicitly redistributed to positive-net-vote Pods through the normalization in the denominator of $s_{j,E}$. Negative voters do not receive emissions; their economic role is to filter low-quality content from the curated dataset.

7.6 Datanet Creation Fee Distribution

When a Datanet is created with fee fee_{create}:

$$\text{Burned} = \beta_{burn} \cdot \text{fee}_{create} \tag{16}$$

$$\text{IncentivizationPool} = \beta_{pool} \cdot \text{fee}_{create} \tag{17}$$

$$\text{Locked} = \beta_{lock} \cdot \text{fee}_{create} \tag{18}$$

The locked portion is returned to the Datanet owner upon burning the Datanet NFT (winding down the Datanet).

7.7 Emission Seeding Tax

When a Datanet owner seeds A tokens of a non-REPPO asset:

$$\text{ToTreasury} = \tau_{tax} \cdot A \tag{19}$$

$$\text{ToEmissionPool} = (1 - \tau_{tax}) \cdot A \tag{20}$$

7.8 Economic Value of Feedback (EVOF)

For a Datanet with N unique voters in epoch E, where voter i has voting power VP_i and casts votes_i votes, and $V = \sum_i \text{votes}_i$ is the total vote count:

$$W_E = \log(N) \cdot \frac{\sum_{i=1}^{N} \sqrt{VP_i} \cdot \text{votes}_i}{V} \tag{21}$$

$$\text{EVOF}_E = W_E^2 \tag{22}$$

Network-level EVOF aggregates across all active datanets:

$$\text{EVOF}_{\text{network},E} = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \text{EVOF}_{d,E} \tag{23}$$

where $\mathcal{D}$ is the set of active datanets in epoch E.

7.9 Equilibrium Conditions

For sustainable protocol operation:

1. The aggregate value captured through burns and taxes should approximate or exceed protocol operating costs:

$$\sum_{\text{datanets}} (\beta_{burn} \cdot \text{fee}_{create} + \tau_{tax} \cdot A_{non\text{-}REPPO}) + \text{PublishingFeeCuts} \geq \text{ProtocolCosts} \tag{24}$$

2. Datanet-level sustainability: a Datanet is viable when its revenue exceeds its emission seeding costs:

$$\text{PublishingFees}_S + \text{AccessFees}_S + \text{ExchangeRevenue}_S > \text{EmissionSeeding}_S \tag{25}$$

3. Token demand equilibrium: as new datanets are created, demand for \$REPPO increases through creation fees and tax-free seeding, while supply decreases through burns:

$$\frac{d(\text{Supply})}{dt} < 0 \quad \text{(net deflationary under normal conditions)} \tag{26}$$

8 Network Effects and Value Accrual

8.1 The Datanet Flywheel

Each Datanet operates a self-reinforcing cycle:

1. The Datanet owner seeds emissions and defines a data collection goal.
2. Publishers submit raw data, paying publishing fees.
3. Voters lock REPPO and evaluate data quality through positive and negative votes.
4. Consensus across multiple epochs produces a curated dataset.
5. The Datanet owner lists the dataset on repo.exchange.
6. Data consumers subscribe, generating recurring revenue.
7. The Datanet owner reinvests a portion of revenue into the emission pool.
8. The cycle repeats with an expanding dataset and growing subscriber base.

This flywheel is sustained as long as the Datanet owner maintains the emission pool. The subscription model on repo.exchange is the key enabler: because datasets update every 48 hours, consumers have reason to maintain ongoing subscriptions rather than making one-time purchases.

8.2 Cross-Datanet Network Effects

As the number of datanets grows, several network effects emerge:

- **Voter liquidity:** A larger pool of veREPPO holders means more voters available for each Datanet, improving the quality and speed of consensus.
- **Data composability:** Downstream consumers can subscribe to multiple datanets and compose datasets across domains—for example, combining street-level imagery data with natural language description data to train multimodal agents.
- **Datanet assessment signals:** As more datanets compete, the performance assessment mechanism (Section 6.4) produces a public ranking of Datanet quality, directing attention and capital toward productive data markets.

8.3 Value Accrual to $REPPO

The $REPPO token accrues value through:

- **Burn pressure:** Every Datanet creation permanently burns ~% of the creation fee.
- **Buy pressure:** Datanet owners who seed emissions in $REPPO avoid the % tax, creating demand for the token on secondary markets.
- **Lock pressure:** Voters must lock $REPPO to participate, reducing circulating supply. The non-linear lock duration bonus incentivizes longer locks.
- **Fee capture:** The protocol's% cut on publishing fees and Data Exchange transaction fees flow to the treasury, available for buybacks or burns at DAO discretion.

8.4 Continuous Data Pipeline as Differentiation

The core value proposition of Reppo is not static dataset creation—there are centralized alternatives for that. The differentiation is the *continuous* nature of the pipeline. Because data is published, evaluated, and curated in 48-hour cycles, downstream consumers receive datasets that track real-world changes in near real time.

For experiential RL applications—robotics systems that need current environmental data, agents that need up-to-date tool documentation, models that need fresh human preference signals—this continuous pipeline is the protocol's core differentiator. A consumer subscribing to a Reppo Datanet is subscribing to a living dataset, not a static file.

9 Token Supply Dynamics

9.1 Token Sources

Reppo does not implement inflationary protocol-level emission mining. The only source of new $REPPO tokens is the initial token supply as defined at genesis. All emissions within datanets are funded by Datanet owners from existing token supply.

9.2 Token Sinks

Multiple mechanisms remove tokens from circulation:

Datanet Creation Burns. ~% of every Datanet creation fee is permanently burned. As the network grows and more datanets are created, cumulative burn increases.

veREPPO Locking. Tokens locked for veREPPO are removed from circulating supply for the duration of the lock. The non-linear voting power bonus ($\gamma > 1$) incentivizes maximum-duration locks, creating sustained illiquidity.

Non-REPPO Emission Tax. When Datanet owners seed emissions in non-REPPO tokens, % flows to the protocol treasury. The DAO can use these diversified treasury assets for REPPO buybacks, further reducing supply.

Protocol Fee Capture. The % cut on publishing fees and Data Exchange transaction fees accumulates in the protocol treasury. The DAO decides allocation: burn, buyback, or operational spending.

9.3 Supply Dynamics

Under normal operating conditions—datanets being created, emissions being seeded, publishers paying fees—the protocol is structurally deflationary:

$$\Delta \text{Supply}_E = -\text{Burns}_E - \text{Buybacks}_E \leq 0 \tag{27}$$

The rate of deflation scales with network activity: more datanets and more publishing volume produce more burns and more treasury inflow available for buybacks.

10 Guardrails and Abuse Mitigation

A capital-weighted data curation system must defend against adversarial behavior at multiple levels: spam publishers, colluding voters, whale manipulation, and sybil attacks. Reppo's architecture embeds several structural guardrails that make gaming economically irrational under normal conditions.

10.1 Economic Barriers to Entry

- **Publishing fees:** Every Pod submission requires a fee set by the Datanet owner. This creates a non-trivial cost floor for spam: an attacker attempting to flood a Datanet with low-quality data must pay per submission, making high-volume spam economically unsustainable.
- **veREPPO lock requirement:** Voters must lock \$REPPO for a chosen duration before they can participate. There is no way to vote without capital at risk. The non-linear lock bonus ($\gamma > 1$) further incentivizes long-term locks, raising the opportunity cost of short-term manipulation.
- **Datanet creation cost:** The 10,000–20,000 \$REPPO creation fee with a% burn makes it expensive to spin up datanets for adversarial purposes.

10.2 Temporal Defenses

- **Linear decay:** Voting power decays linearly over the 48-hour epoch (Section 4.2). Last-minute vote manipulation—where an attacker waits to observe consensus and then piles in—is penalized by reduced effective power. A vote cast at the midpoint carries half the weight of one cast at the start.
- **Epoch-based settlement:** All outcomes resolve at epoch boundaries, not continuously. This prevents front-running and ensures that all voters face the same information horizon within an epoch.
- **Multi-epoch convergence:** The synthetic oracle (Section 4.2.4) relies on repeated consensus over six to seven epochs. Manipulating a single epoch has limited impact on the cumulative quality signal; sustained manipulation across many epochs is prohibitively expensive.

10.3 Whale and Collusion Resistance

- **Square root dampening in EVOF:** The term $\sqrt{VP_i}$ in the EVOF metric (Section 6.5) ensures that large stakeholders have a diminishing marginal influence on the confidence signal. A whale with $4\times$ the voting power of a smaller voter contributes only $2\times$ the signal weight.
- **Net-vote mechanism:** Negative voting provides an explicit channel for dissent. If a whale attempts to push a low-quality Pod into the curated set, smaller voters can collectively vote against it. Because emissions for negative-net-vote Pods are zeroed out and redistributed, there is a direct economic incentive for honest voters to challenge manipulated outcomes.
- **Voter diversity through EVOF:** The $\log(N)$ term in EVOF rewards broad participation. A Datanet that relies on a small number of whale voters will have a lower EVOF than one with diverse participation at comparable total stake, creating a visible signal that the curation may be less robust.

10.4 Sybil Resistance

Sybil attacks – where one entity creates many accounts to amplify influence – are mitigated by the capital requirement: splitting stake across k sybil accounts does not increase total voting power (since VP is a function of locked tokens, not account count). The $\log(N)$ diversity bonus in EVOF could in theory be inflated by sybils, but because each sybil account must independently lock capital, the cost of meaningfully inflating N is linear in the capital deployed.

For additional identity-layer defense, the protocol plans to integrate with proof-of-personhood systems to distinguish human voters from AI agents, adding a complementary non-economic sybil resistance layer.

10.5 Data Quality Feedback Loops

Beyond structural defenses, the EVOF metric (Section 6.5) and the Datanet performance assessment mechanism (Section 6.4) create continuous feedback loops:

- Datanet owners monitor their EVOF trajectory against the network benchmark. A declining EVOF relative to peers signals potential quality degradation or gaming, prompting corrective action (adjusting fees, tightening publisher guidelines, or increasing emissions to attract more honest voters).
- veREPPO holders assessing Datanet performance have an economic incentive to identify and avoid datanets with suspicious voting patterns, directing capital toward better-run data markets.

11 Roadmap

- **V1 (Live):** datanets with stake-assured prediction markets, commit-reveal voting, centrally managed publishing fees, Data Exchange beta (repo.exchange).
- **V2 (Q2 2026):** Self-served Datanet creation, Datanet NFTs, net voting (positive/negative), linear decay, removal of commit-reveal, configurable publisher/voter splits, non-REPPO emission seeding with tax, third-party incentives mechanism, Datanet performance assessment.
- **V3 (Future):** Cross-Datanet data markets, enterprise-grade access controls, composable RL task environments, Infrastructure Exchange integration.

12 Conclusion

Reppo provides tokenized infrastructure for the production of continuously curated datasets for experiential reinforcement learning. The protocol's design is grounded in a simple economic primitive: datanets as on-chain data businesses where publishers contribute raw data, voters lock capital behind quality assessments, and the resulting consensus produces a synthetic oracle that requires no external arbiter.

The V2 architecture introduces net voting with linear decay, self-served Datanet creation with NFT-based ownership, a non-inflationary emission model where Datanet owners fund their own incentive pools, and the Economic Value of Feedback (EVOF) metric that gives Datanet owners and dataset consumers a live confidence signal for dataset quality. These mechanisms, combined with structural guardrails against spam, collusion, and whale manipulation, produce a robust and structurally deflationary token economy tied directly to data production activity.

By organizing data curation as a prediction market and data consumption as a subscription service, Reppo creates a continuous pipeline from raw human experience to curated RL training data–the infrastructure layer that experiential AI systems require to learn from the real world.

References

1. Silver, D., & Sutton, R. S. (2025). *Welcome to the Era of Experience*. Retrieved from DeepMind.

2. Curve Finance. (2020). *Curve DAO: Vote-escrowed CRV and veTokenomics*. Retrieved from https://resources.curve.fi/base-features/understanding-crv

3. Aerodrome Finance. (2023). *Aerodrome: The central trading and liquidity marketplace on Base*. Retrieved from https://aerodrome.finance/docs