Panoramic Governance

0xCygaar, 0xBeans, 0xMons

October 6, 2024

Abstract

We introduce Panoramic Governance, a new mechanism that incentivizes governance participation and protocol growth within Layer 2 blockchains. Panoramic Governance can be implemented on any Layer 2 blockchain with a native governance token and a sequencing system where users pay gas fees for faster transaction ordering and trusted pre-confirmations. PG consists of two interdependent systems: one that distributes sequencer fees to governance voters that actively participate in the network and one that allows voters to direct token emissions to protocols built on the given blockchain. Protocols increase demand for block space, resulting in higher sequencer fees. This, in turn, boosts earnings for active governance voters, who are then motivated to allocate more emissions to successful protocols on the chain.

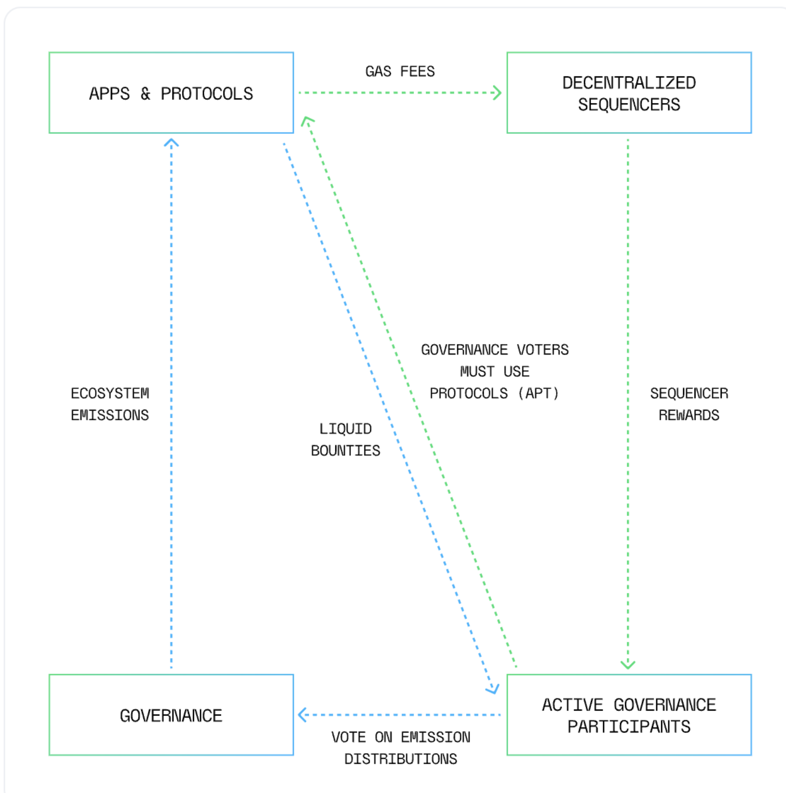

1 Introduction

Rollups, a type of Layer 2 blockchain, have exploded in popularity since late 2022 and are now the primary scaling solution for Ethereum [4], offering lower fees and higher throughput compared to mainnet. Since October 2022, Ethereum rollups have consistently processed more transactions per second than Ethereum mainnet. As of the time of this writing, Ethereum-based rollup activity is nearly ten times greater than that of the Ethereum blockchain [1]. For the purposes of this paper, “rollups” will be defined as Ethereum-based rollups that use Ether as the native gas token.

Most rollups today utilize sequencers - systems that accept, order, and publish transactions in batches to Ethereum (or alternate data availability layers) which then provides a final ordering of the transactions. Transactions processed by sequencers require gas fees, similar to standard Layer 1 blockchains. These gas fees are composed of two parts: one part to pay the cost of posting calldata to the data availability layer, and one part to account for state modifications and computations done on the Layer 2 virtual machine itself.

The compute and state-modifying portions of transaction fees are typically accrued directly to the rollup operator. This differs from Layer 1 blockchains, where gas fees are typically burned as a deflationary measure. In contrast, most rollups use Ether as the native gas currency, eliminating the need for either deflationary or inflationary mechanisms because the token does not originate on the rollup itself. As a result, rollup operators end up accumulating Ether, oftentimes in large quantities [2][3][9]. The more gas that is spent onchain, the more fees a rollup operator accrues. These fees have historically been accrued by the rollup operator or kept in a treasury to be allocated according to governance mandates at some point in the future.

This paper seeks to introduce a new mechanism where the governance voters, active users, and applications of a rollup become the primary beneficiaries of the chain's success, rather than the operator itself. As all participants in the rollup (users, protocols, voters) play an integral part in growing the network and making sure it is as successful as possible, this paper posits that they should be the ones benefiting from that success.

2 Sequencer Fees

In this section, we provide an overview of sequencer fees, explaining how they are accrued and the various factors that influence them. L2s primarily collect transaction (gas) fees to cover the cost of operating sequencers and to adjust costs based on blockspace demand. This differs from L1s which use gas fees to pay validators and to burn supply as a deflationary measure.

In the hypothetical rollup we discuss below, we assume a standard Ethereum Virtual Machine (EVM) setup [5] with all fees denominated in Ether rather than a separate gas token.

2.1 Transaction Costs

In order to formally model the sequencer fees earned by an L2 blockchain, we first need to define how the costs of an individual transaction are computed. Gas is the single unit of measurement for the usage of compute, storage access, and calldata in EVM transactions. Gas fees are used to price the cost of each gas unit for a given block.

For an L2 transaction, costs are broken down into the $L1_{Fee}$, the cost to submit the compressed transaction calldata to the data availability layer, and the $L2_{Fee}$, the cost of executing the transaction on the L2's EVM. The $L1_{Fee}$ is used purely for posting calldata or statediffs to the DA layer, so they are not taken into account when computing sequencer revenue.

The $L2_{Fee}$ can be expressed with the following formula, where F_B is the base fee, F_P is the priority fee, and $L2_{Gas\_Used}$ is the amount of gas consumed for L2 execution:

$$L2_{Fee} = L2_{Gas\_Used} \times (F_B + F_P)$$

The base fee is set per block, can fluctuate based on activity levels on the chain (some chains may keep the base fee static), and sits at a floor price based on the rollup's configuration. The formula used to compute the base fee of an L2 varies per chain, but in general the more demand there is for blockspace within a set period of time, the higher the base fee goes.

Priority fees are used to prioritize transactions within a block. The higher the priority fee, the earlier a transaction is executed within a block. However, since some rollups [7] do not use priority fees and those that do show only minimal impacts, this paper omits them from our calculations for simplicity.

2.2 Total Value Accrual

We can then represent the total rewards accrued per block as such (assuming the base fee per block is consistent):

$$R_{block} = \sum_{i=0}^{n} L2_{Fee_i} \tag{1}$$

We can then represent the total rewards accrued for a given period of time as:

$$R_t = \sum_{j=0}^{m} R_{block_j} \tag{2}$$

This shows us that the total rewards accrued by a rollup's sequencers is a function of the total amount of execution gas consumed by transactions conducted on the network. In general, this means the more activity there is on the chain, the larger the reward accumulation for a rollup will be. Compare this to Ethereum (or other similar L1 blockchains) where higher activity on chain results in gas fees being burnt rather than direct value accrual to validators. L2s are unique in their ability to capture execution fees; PG aims to leverage these fees to reward active governance participants.

2.3 Decentralized Sequencing

Although most rollups today use a centralized sequencer, PG works best as part of a rollup with a decentralized sequencer. In this system, multiple sequencers take turns sequencing each block and reach a consensus to determine the ordering and validity of transactions. The rest of this paper assumes that the rollup contains a decentralized sequencer.

In most decentralized sequencing designs, sequencers earn revenue through a combination of sequencer fees (as described above) and Maximum Extractable Value ("MEV") rewards. In PG, sequencers will retain a small portion of the sequencer fees for themselves, and direct the rest of the fees to a vault contract that can later distribute fees autonomously to active governance participants. Also, a subsidy can be given to sequencer operators to incentivize participation in the sequencing network.

3 Governance Alignment

Governance is critical to any decentralized network, and is an important pillar of PG. Governance token holders are the recipients of sequencer fees and decide where to direct ecosystem emissions. However, without any mechanism for enforcing alignment, a malicious actor can quickly acquire a large number of governance tokens, sway a vote towards a harmful proposal, and divest their tokens immediately after.

One potential solution to this problem is to require tokens to be "vote-escrowed" where governance tokens are deposited into a smart contract and locked for a duration of time. In return for escrowing tokens, token holders will receive an equivalent number of transfer-locked vote-escrowed tokens, representing their locked positions. Vote-escrowed token owners will then be the recipients of the sequencer fees. Additionally, the longer a token holder escrows their tokens, the more governance power they receive. By time-locking their governance tokens, token holders are committing to the long term success of the chain. While this does not entirely eliminate the risk of governance capture mentioned earlier, it does increase the cost of such actions, creating a deterrent effect. This time-locked approach is heavily inspired by Curve Finance's governance system [6].

There are several other ways to reduce the risk of governance attacks including requiring voters to go through Know Your Customer ("KYC") protocols, confirming each voter's reputation score, or having strong veto powers in place [10]. However, we believe that requiring tokens to be escrowed in exchange for voting powers strikes the right balance between preserving decentralization, while also correctly aligning incentives and promoting efficient and engaged governance.

4 Active Participation Threshold

Active Participation Threshold ("APT") is a mechanism designed to ensure that only users who actively contribute to the network receive rewards. This mechanism helps maintain network health, encourages organic activity, and deters passive or exploitative behaviors. APT establishes a system where user rewards are scaled based on an individual's network activity relative to the activity of others. In essence, rewards are only given to users that are actively using the network.

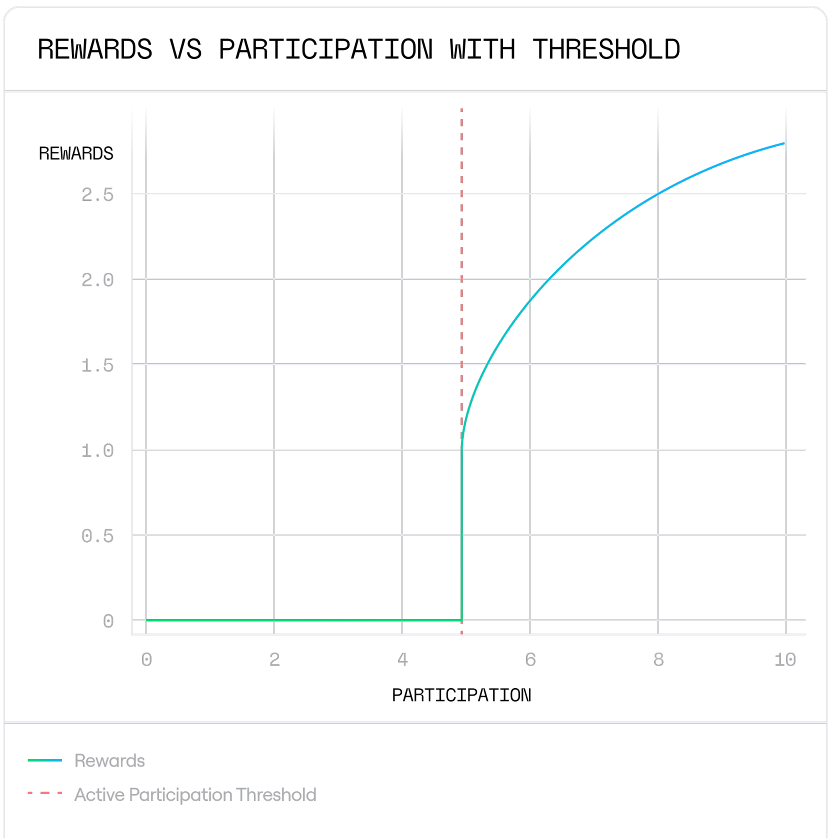

The criteria for what dictates an "active" participant for a network varies and should be defined by governance. Metrics such as transaction count, gas spent, total value transacted onchain, or time active can be used individually or together to determine if a user has been "active." For the scope of this paper, we will use transaction count as the primary metric for determining activity as it's simple to compute, can be easily determined through available onchain data, and serves as a relatively close approximation for total gas spent onchain. Implementations of PG can use a combination of the metrics listed above.

Below, we propose a formula that rewards governance participants only if they meet the necessary transaction count threshold. We want the APT to be a dynamic value directly correlated to network growth and reward distribution. In general, the higher the median user transaction count is for an epoch, the more fees the sequencers have accumulated. As the network grows, rewards will increase for token escrowers in subsequent epochs. A naive approach would be:

$$user\_threshold = \frac{median(tx\_count) \times users\_escrowed\_amt}{avg\_escrowed\_amt} \tag{3}$$

where:

- median(tx_count) is the median transaction count of every token escrower for this epoch.
- avg_escrowed_amt is the average tokens escrowed for this user.

This method determines a user's activity threshold by scaling it according to the amount they have escrowed relative to other token holders. However, this method is vulnerable to griefing attacks where malicious actors spam the chain with low value transactions to increase the median transaction count in order to increase the overall activity threshold. To prevent this from happening, we can use a weighted transaction fee value approach.

4.1 Weighted Median Transaction Count

The median transaction count, weighted by average transaction fee value, for the current epoch is defined as:

1. Find each users average transaction value:

$$avg\_transaction\_value_i = \frac{\sum_{j=1}^{total\_txn\_count_i} txn\_value_{ij}}{total\_txn\_count_i} \tag{4}$$

2. Sort user transaction counts $(\text{total\_txn\_count}_1, \text{avg\_txn\_value}_1), \ldots, (\text{total\_txn\_count}_n, \text{avg\_txn\_value}_n)$ where $\text{avg\_txn\_value}_1 \leq \ldots \leq \text{avg\_txn\_value}_n$.
3. Find the median txn_count in the above sorted list.

Using the weighted median txn count, we can calculate a users threshold using this formula:

$$user\_threshold = \frac{weighted\_median(txn\_count) \times users\_escrowed\_amt}{avg\_escrowed\_amt} \tag{5}$$

Using a weighted transaction count can significantly reduce the likelihood of a griefing attack. By assigning greater weights to higher-value transactions, it becomes economically impractical for malicious actors to flood the network with low-value transactions aimed at manipulating the weighted median transaction count. The effectiveness of the method increases as the network expands and more legitimate users conduct onchain transactions.

If a malicious actor tries to manipulate the median value by spamming the chain with higher-value transactions, the ecosystem and its participants will benefit from an increase in sequencer rewards. However, it's important to note that while adopting this method reduces some of the risks posed by malicious actors, it does not entirely eliminate them. There should be fail-safe protective mechanisms in place that governance participants can activate in the early stages of the network.

While the above is 1 example, every network should implement its own APT heuristic that best protects the network from attacks.

5 Activity Driven Loop

The first module of PG is the activity driven feedback loop. There are three components involved in this loop: protocols, sequencers, and governance participants.

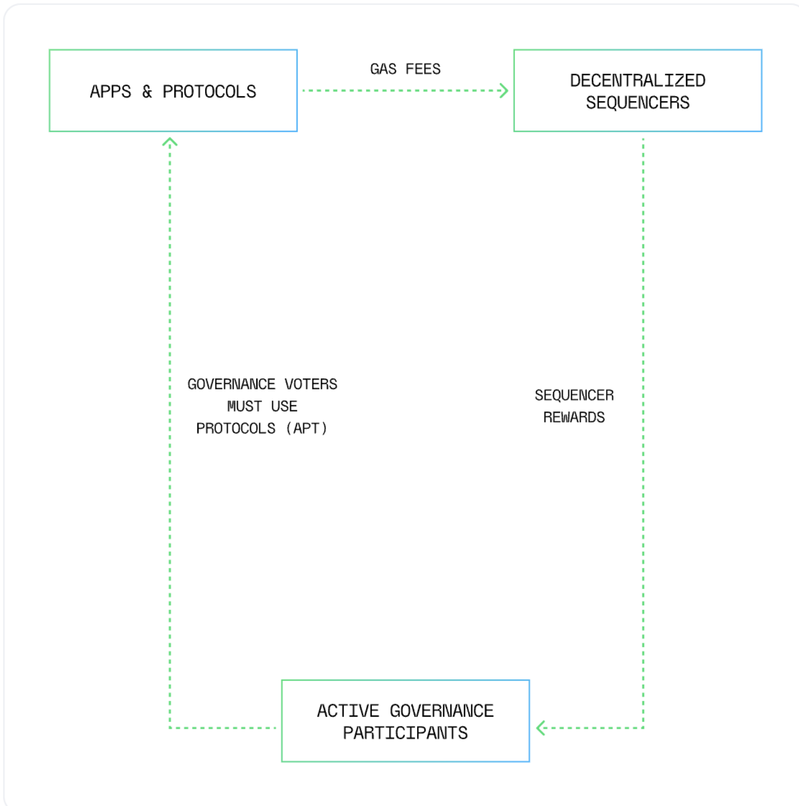

5.1 Decentralized Apps / Protocols

These are the applications or protocols deployed on the chain. They are composed of one or more smart contracts and require some form of on-chain functionality. Examples of decentralized apps or protocols include decentralized exchanges, onchain social networks, onchain games, lending markets, NFT projects, and more. When deployed on the chain, these applications or protocols drive blockspace usage and demand which leads to gas consumption. Generally, the more onchain activity an application or protocol has, the more sequencer fees it delivers to the network.

5.2 Sequencers

As discussed previously, sequencers accrue value whenever a transaction is submitted and executed. With enough activity and a reasonably-metered base fee, sequencers can accrue a significant amount of fees. In our proposed activity driven loop, these fees can be redirected to active governance voters rather than being kept as revenue for the sequencer operator. Therefore, the more gas consumed on the chain, the more fees are directed towards active chain participants.

5.3 Active Governance Participants

Governance participants are the direct beneficiaries of increased network activity as they receive the chain's sequencer fees. This setup is vastly different from that of traditional rollups, as the primary beneficiaries of those systems were the rollup operators themselves. Governance voters create a social layer which, while not directly measurable, benefits the ecosystem by encouraging the development of more applications on the chain and supporting the success of existing protocols.

In order to receive a portion of the sequencer rewards, governance token holders must:

- Escrow their tokens for a specified period of time, in order to participate in governance;
- Participate in the voting of ecosystem emissions (described later on); and
- Be active participants of the network as determined by the Active Participation Threshold.

In essence, token holders cannot passively earn rewards within a PG rollup - they must actively participate in network governance and be recurring users of the chain itself to be rewarded.

5.4 Positive Feedback Loop

All three components work together to promote the growth of the network. Onchain applications generate blockspace demand by creating applications and flows that users want to use. The network's sequencers act as a value capture layer for the gas consumed onchain. Governance voters that actively use the network receive rewards based on the total sequencer fees accrued. They are then motivated to push more applications to build onchain, which increases demand for block space, boosts gas consumption, and further enhances the overall sequencer rewards earned.

6 Ecosystem Emissions

Emissions can be used as a mechanism to incentivize ecosystem development. Inspired by Curve Finance's gauge mechanism [6], we propose an approach that allocates a percentage of tokens for distribution among various protocols. Holders of vote-escrowed tokens with governance powers vote to determine the allocation of these emissions as rewards each epoch. There is one voting period per epoch, and there should be a predetermined number of epochs at the genesis of the Panoramic Governance system. These emissions should be leveraged to grant successful protocols a greater stake in the network's governance and thus a larger share of the generated sequencer fees.

6.1 Emissions Rate

Emission rewards are granted to protocols per epoch based on the following formula:

$$r_p = \frac{v_p}{V_t} \times R \tag{6}$$

Where:

- r_p is the rewards emitted to protocol p
- v_p is votes given to protocol p
- V_t is total votes
- R is total rewards per epoch

During a single epoch, a protocol will receive a pro rata share of the epoch's rewards based on the percentage of total votes it receives. It is up to each implementation to determine the R value per epoch. This amount should account for the total duration of the ecosystem emissions process, the number of epochs it spans, and the total emission rewards allocation set at genesis.

The total emissions of all ecosystem rewards can be represented by the following formula:

$$\sum_{i=1}^{n}\left(\frac{v_{p_i}}{V_{t_i}} \times R_i\right) = X \tag{7}$$

Where:

- i is the current rewards epoch
- v_{p_i} is votes given to protocol p for epoch i
- V_{t_i} is total votes for epoch i
- R_i is total rewards for epoch i
- X is the total emissions set

X should be consistent with the set amount of emissions that were initially allocated for ecosystem rewards.

Note: The formulas above assume that there is at least one voter in the system. Implementations can decide what to do when there are no voters - options include pausing emissions or using other heuristics to determine distributions.

6.2 Identifying Eligible Protocols

A fully permissionless voting system can lead to a malicious actor acquiring a large portion of the native governance token, escrowing it, and voting to direct emissions towards themselves. While the lock-to-vote mechanism helps deter this attack vector, a malicious attack is still viable if there are large enough incentives to do so.

To prevent such an attack, especially early on when the network is not yet fully decentralized, we recommend using a permissioned system where only certain protocols are eligible for emissions. This eligibility can be determined based on objective metrics or through subjective heuristics shaped by social consensus. Over time, when the network is more decentralized, the system can be made permissionless, allowing any protocol to be eligible for rewards.

6.3 Liquid Bounties

With the token emission mechanism outlined above, we foresee the emergence of liquid bounty markets - a system where protocols offer incentive bounties to token holders in exchange for emissions votes. This acts as a secondary reward mechanism for holders that escrow their tokens, and helps protocols align with the largest and most active participants in the network. For protocols to offer sustainable incentive bounties over the long term, they must consistently accrue and generate value.

7 Emissions Driven Loop

The second module of PG is the emissions driven feedback loop. There are three components involved in this loop: ecosystem emissions, governance participants, and protocols.

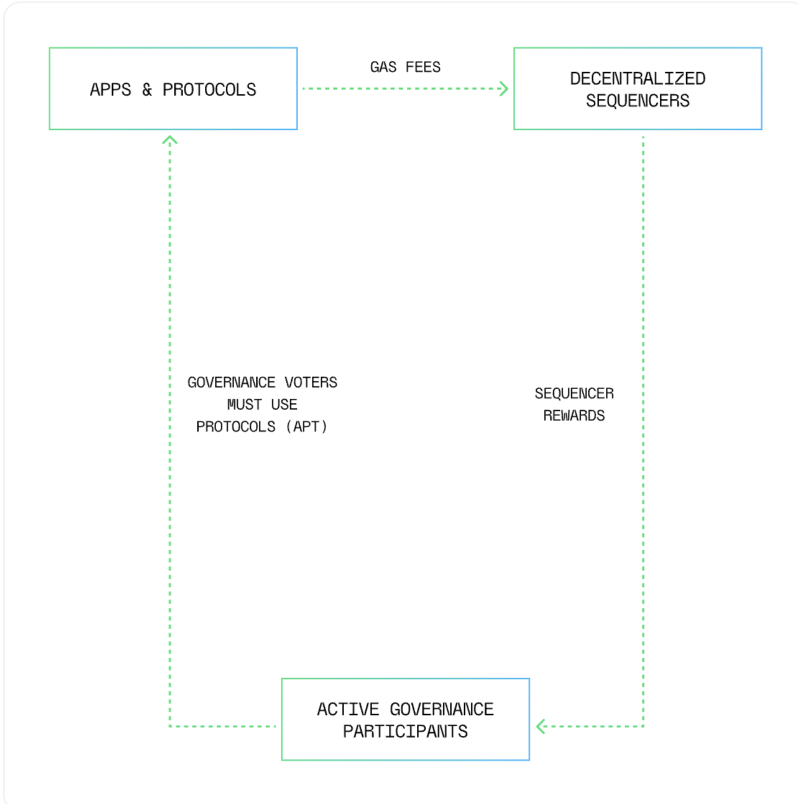

7.1 Ecosystem Emissions

These emissions are used to reward protocols that are deployed and add value to the network. They are emitted every epoch, with the total duration and number of tokens emitted determined ahead of time. Governance participants determine which protocols receive the emissions.

7.2 Governance Participants

Token holders will vote on which protocols to send token emissions to through an onchain governance system. Although methods for ensuring governance alignment may differ, generally, owning more governance tokens should result in greater voting weight. Voters will be incentivized to vote for the protocols that bring the most value to the network and the protocols that offer the largest liquid bounties.

7.3 Onchain Protocols

Protocols deployed on the network are the recipients of the pre-allocated ecosystem emissions. Among eligible protocols, those that bring the most value and activity to the network should receive the largest percentage of the emissions. Examples of beneficial apps for the ecosystem include SocialFi apps that bring millions of non-crypto users to the chain, decentralized exchanges with high trading volumes driving significant chain usage, and bridging protocols that attract substantial Total Value Locked ("TVL") into the chain. In each example, the apps are either bringing more users, transactions, or capital to the chain, leading to overall ecosystem growth. Ecosystem emissions grant beneficial protocols a larger stake in the network, enabling them to share in the success of the chain.

7.4 Positive Feedback Loop

Together, emissions, governance participants, and protocols create a dynamic system that promotes the growth of the network. Emissions are used to reward the protocols that drive the most activity to the chain, governance participants vote on how to allocate the emissions, and protocols continue to improve their products to further increase network activity.

8 Conclusion

By using the PG system described above, we can create a dual module system that rewards active network and governance participation while also incentivizing ecosystem growth. Each module can function independently to drive network adoption, but they also work together in a symbiotic fashion to align token holders, onchain protocols, and application users to grow the ecosystem sustainably and robustly.

The dual module system operates in a closed environment without relying on artificial token inflation. All rewards are sourced directly from a predefined tranche of tokens and capital that is injected into the system through sequencer fees.

9 Sources

1. https://l2beat.com/scaling/activity
2. https://dune.com/jhackworth/base-metrics
3. https://defillama.com/fees/zksync-era
4. https://ethereum.org/en/roadmap/scaling/
5. Page 28, https://ethereum.github.io/yellowpaper/paper.pdf
6. https://classic.curve.fi/files/CurveDAO.pdf
7. https://docs.arbitrum.io/learn-more/faq#do-i-need-to-pay-a-tip-or-priority-fee-for-my-arbitrum-transactions
8. https://studio.glassnode.com/metrics?a=ETH&m=fees.GasUsedMean&s=1704085200&u=1714708799&zoom=
9. https://dune.com/niftytable/rollup-economics
10. https://a16zcrypto.com/posts/article/dao-governance-attacks-and-how-to-avoid-them/