Information Systems 139 (2026) 102695

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Measuring the decentralisation of DeFi development: An empirical analysis of contributor distribution in Lido☆

Giuseppe Destefanis [a],*, Jiahua Xu [a,b,c], Silvia Bartolucci [a]

[a] Department of Computer Science, University College London, London, United Kingdom
[b] Centre for Blockchain Technologies, University College London, London, United Kingdom
[c] DLT Science Foundation, London, United Kingdom

ARTICLE INFO

Keywords:
DeFi
Contributor structure
Decentralisation

ABSTRACT

Decentralised finance (DeFi) protocols often claim to implement decentralised governance via mechanisms such as decentralised autonomous organisations (DAOs), yet the structure of their development processes is rarely examined in detail. This study presents an in-depth case analysis of the development activity distribution in Lido, a prominent DeFi liquid staking protocol. We analyse 6741 human-generated GitHub actions recorded from September 2020 to February 2025. Using standard inequality metrics – Gini coefficient and Herfindahl–Hirschman Index – alongside contributors' interaction network and core–periphery modelling, we find that development activity is highly concentrated. Overall, the weighted Gini coefficient reaches 0.82 and the most active contributor alone accounts for 24% of the total activity. Despite an even split between core and peripheral contributors, the core group accounts for 98.1% of all weighted development actions. The temporal analysis shows an increase in concentration over time, with the Gini coefficient rising from 0.686 in the bootstrap phase to 0.817 in the maturity phase. The contributors' interaction network analysis reveals a hub-and-spoke structure with high centralisation in communication flows. While a case study of a single protocol, Lido represents a critical test of decentralisation claims given its prominence, maturity, and DAO governance structure. These findings demonstrate that open-source DeFi development can exhibit highly concentrated control patterns despite decentralised governance mechanisms, revealing a persistent gap between governance and operational decentralisation.

1. Introduction

Over the past few years, decentralised finance (DeFi) has evolved from a niche innovation into a major segment of the global financial ecosystem. Built on public blockchains and powered by smart contracts, DeFi protocols offer financial services — such as lending, trading, staking, and asset management — without the need for traditional intermediaries [1–3].

Since 2020, the sector has experienced exponential growth, with total value locked (TVL) in DeFi protocols peaking at over $180 billion. This surge has been driven by user adoption, venture capital investment, and the proliferation of decentralised applications (dApps) across multiple blockchain networks [4]. DeFi has become a core component of the broader Web3 movement, attracting both retail and institutional participants seeking open, permissionless alternatives to traditional finance [5].

However, the rapid expansion of DeFi has also introduced substantial risks. Security breaches, smart contract vulnerabilities, governance attacks, and fraud have exposed weaknesses in protocol design and operations [6–10]. The composability of DeFi — where protocols interconnect and depend on one another — may amplify systemic risks [11], prompting regulatory concerns over market integrity and financial stability [12–14].

Beyond technical and financial risks, questions have emerged around the true extent of decentralisation in DeFi. Despite decentralised governance mechanisms such as DAOs, power often remains concentrated in areas such as governance token distribution, validators' control, and development authority [15,16]. This "decentralisation illusion" [17] reflects the reliance of many DeFi systems on small groups of actors for essential functions such as protocol upgrades, governance coordination, and maintenance.

Open-source development is central to most DeFi protocols, with platforms like GitHub serving as hubs for collaboration, issue tracking, and code contributions [11,18–20]. In theory, this model promotes

☆ This article is part of a Special issue entitled: 'BISs' published in Information Systems.

* Corresponding author.
E-mail address: g.destefanis@ucl.ac.uk (G. Destefanis).

https://doi.org/10.1016/j.is.2026.102695
Received 2 August 2025; Received in revised form 14 November 2025; Accepted 27 January 2026
Available online 29 January 2026

transparency and decentralised control. Public access to source code allows diverse contributors to participate in auditing, feature development, and security review. Yet, open-source status alone does not ensure decentralisation. In practice, control over repositories, authorship of key components, and decision-making authority may be restricted to a few core contributors or affiliated organisations, creating potential governance bottlenecks and resilience challenges.

Understanding the distribution of development effort is therefore essential. Projects with a high concentration of contributors may face risks related to governance capture, single-contributor dependency, and maintenance bottlenecks. Conversely, broader participation can enhance resilience by distributing auditing, innovation, and maintenance responsibilities across a wider base. Using publicly accessible GitHub data, it is possible to empirically assess how development activities are distributed, how collaboration patterns evolve, and whether core development tasks — such as pull request review or issue resolution — are concentrated among a limited subset of individuals.

This study focuses on Lido, one of the most prominent DeFi protocols, as a case for analysing development decentralisation. Lido provides liquid staking solutions across multiple blockchain networks, allowing users to stake assets such as ETH, SOL, MATIC, and DOT while maintaining liquidity through "derivative tokens"[1] like stETH [22]. Governed by the Lido DAO, the protocol is designed to reflect decentralised governance principles. However, as with many DeFi projects, critical areas — including smart contract control, validator selection, and core development — remain partially centralised.

Lido's repositories are hosted on GitHub [23], theoretically allowing open participation. Yet, the contribution landscape is dominated by a core development team, raising concerns about contributor diversity and resilience. Lido serves as a compelling case study for several reasons. It is the largest liquid staking protocol by TVL, with over $30 billion at peak, making its development practices systemically important. Its 54-month GitHub history provides an important longitudinal dataset capturing governance changes, technical upgrades, and market cycles. It also exemplifies the decentralisation paradox: fully open-source and DAO-governed, yet subject to technical and coordination constraints that promote centralisation. Finally, its development records enable full empirical observation, avoiding the selection biases seen in projects with fragmented repositories.

This study employs a single-case exploratory design to evaluate methods for measuring development concentration in DeFi. Following Yin's [24] framework, Lido represents both a *critical case* for testing assumptions about decentralisation in open-source DeFi development, and a *revelatory case* providing complete longitudinal data typically unavailable for protocol analysis. While not aiming for generalisation across the entire ecosystem, this work establishes a methodological basis for future comparative work.

We extract structured GitHub activity records, including commits, pull requests, issues, and comments, for Lido, and group them by contributor and calendar month. To quantify the distribution of development activity, we apply inequality metrics (Gini coefficient [25], Herfindahl–Hirschman Index (HHI) [26]), top-N contributor shares, and interaction network centrality scores based on co-participation in discussions.

We address the following research questions:

RQ1 How is development activity distributed among contributors in Lido?
RQ2 How do contribution patterns and collaboration structures evolve over time?
RQ3 How do external market shocks and governance decisions affect the distribution and temporal patterns of development activity in Lido?

[1] A "derivative token" has its value derived from its underlying [21].

RQ1 establishes whether development is broadly distributed or concentrated. **RQ2** introduces a temporal dimension to examine trends in centralisation and analyses comment-based contributor networks using centrality metrics to capture structural collaboration dynamics. **RQ3** examines whether external events reinforce or disrupt existing contribution patterns.

By integrating activity-based and interaction-based analyses, this study offers a detailed empirical evaluation of open-source development in DeFi. The results have implications for assessing decentralisation claims, understanding governance structures, and identifying sustainability risks in protocols' development.

2. Related work

Our investigation of centralisation of development in DeFi protocols draws on several research areas. We examine existing work on the decentralisation illusion in DeFi, concentration patterns in open-source development, and core–periphery network structures.

2.1. The decentralisation illusion in DeFi

[17] introduced the concept of the "decentralisation illusion" in their analysis of DeFi risks, arguing that "full decentralisation in DeFi is illusory" because "algorithm incompleteness" makes some level of centralisation inevitable. They state that "all DeFi platforms have an element of centralisation, which typically revolves around holders of 'governance tokens' (often platform developers) who vote on proposals, not unlike corporate shareholders." While their work provides theoretical grounding for why DeFi systems tend towards centralisation, they do not empirically measure development activity patterns.

In a related vein, [27] examined how changes in the U.S. federal funds rate affect the DeFi sector, finding "a statistically significant and economically important counter-cyclical effect on DeFi lending rates and DeFi asset growth." This empirical evidence challenges the common assumption that DeFi operates independently of traditional financial systems, and instead suggests that monetary policy in centralised finance exerts a measurable influence on DeFi activity. The authors describe this phenomenon as a "decentralisation illusion", highlighting the disconnect between DeFi's decentralised technical architecture and its economic dependence on macro-financial conditions. This work focuses on financial performance and investor behaviour rather than on the development infrastructure underpinning DeFi protocols.

2.2. Concentration in open-source software development

[28] conducted an analysis of 263 Apache Software Foundation projects, finding a severe concentration in developer contributions. They report: "Among the 263 analysed cases, 100 (38.02%) cases are in the range of 0.7–0.8, while 234 (88.97%) of the analysed population is between 0.6 and 0.9" for Gini coefficient values. Their study "undermines the widespread idealistic belief that open-source development is a wide collaborative movement," showing instead that projects were "created by a small, but very active, group of individual, separate contributors."

[29] proposed using the Gini coefficient to measure team activities in open-source software development, specifically noting its application to GitHub projects. Although their work establishes methodological precedent for using inequality metrics in software development contexts, detailed findings about temporal patterns or specific concentration levels were not presented.

Beyond Apache projects, similar concentration patterns have been observed across diverse open-source ecosystems. [30] studied 661 public and 171 enterprise GitHub projects, finding that the "hero" pattern, where 20% of contributors produce 80% of work, is prevalent across both OSS and commercial software development. [31] provided early evidence of this concentration in Apache and Mozilla projects, establishing foundational benchmarks for understanding OSS contribution patterns. More recent work [32] examined the Solidity smart contract language over a 10-year period and found that the top 1% of contributors accounted for 85% of weighted activity, demonstrating that blockchain-related projects exhibit particularly high concentration. [33] further demonstrated that concentrated leadership structures are characteristic of successful software development teams. These findings establish that contributor inequality is endemic to open-source development, with Gini coefficients typically ranging from 0.6 to 0.9 and top contributor shares often exceeding 70%–80% of total activity.

2.3. Core-periphery network structures

The core–periphery paradigm has been extensively studied in network science. [34] developed methods to investigate "the meso-scale feature known as core–periphery structure, which entails identifying densely connected core nodes and sparsely connected peripheral nodes." Applying this concept to software development, [35] studied the impact of so-called different core–periphery movements on Open Source projects, concluding that "a steady core–periphery shift towards the core is beneficial to the project." Their work established the relevance of core–periphery analysis for understanding software development dynamics, though they did not examine blockchain or DeFi projects specifically.

2.4. Blockchain development activity

Limited empirical research exists on development patterns in blockchain projects. CryptoMiso[2] provides rankings of cryptocurrencies based on GitHub commit history, offering descriptive statistics but no analysis of contribution distributions. General observations about blockchain development suggest significant activity, for instance, various sources report thousands of active blockchain developers monthly, but a systematic analysis of concentration patterns remains absent from the literature.

2.5. Research gap

Existing literature has established that (1) DeFi suffers from a decentralisation illusion in governance, (2) traditional open-source projects exhibit high contributor concentration with Gini coefficients typically between 0.6 and 0.9, (3) core–periphery structures characterise collaborative networks. However, no prior work has empirically measured development concentration specifically in DeFi protocols

Our study addresses this gap by applying established concentration metrics and network analysis techniques to examine whether DeFi's unique characteristics, including token incentives and decentralised autonomous organisation (DAO) governance, may alter the development patterns observed in traditional open-source projects. By analysing 54 months of development data from Lido, we provide the first large empirical assessment of the "decentralisation illusion" as it manifests in actual development activities, rather than in theoretical governance structures.

3. Dataset & methods

For this study, we collected data from the core implementation repositories of the Lido Finance GitHub project [23], with a focus on protocol-level smart contracts and essential infrastructure. The dataset consists of structured CSV files that capture key elements of development and collaboration. These include source code modifications (commits), proposed changes (pull requests), issue reporting and enhancement suggestions (issues), and related discussions (comments). Each record contains a contributor identifier and a timestamp, enabling both *temporal* and *contributor*-level analysis.

The following files were extracted for analysis:

- `commits.csv` Records of code commits with author identifiers and timestamps.
- `file_commits.csv` Links commits to modified files, supporting analysis of codebase evolution.
- `files.csv` Metadata on tracked files (e.g., filenames, paths).
- `pull_requests.csv` Information on submitted pull requests, including descriptions, status, and timestamps.
- `issues.csv` Reports of bugs, feature requests, or improvement suggestions, along with associated lifecycle metadata.
- `comments.csv` Contributor discussions on issues and pull requests, used to model interaction networks.

This information was extracted for the entire lifetime of the Lido project, from September 2020 to February 2025.

GitHub repositories often include also automated bot accounts that perform routine maintenance tasks such as dependency updates, test notifications, code formatting, and continuous integration processes. These accounts can distort development metrics by introducing activity patterns that do not reflect human engagement. To address this, we identified and excluded bot accounts using several criteria: account naming conventions (e.g., containing terms such as `bot`, `automated`, or suffixes like `[bot]`), activity patterns indicative of automation (e.g., highly regular commit intervals or exclusively automated pull request creation), and GitHub's built-in bot designation features where available. This preprocessing step was necessary to ensure that the analysis of contribution patterns reflects human contributor behaviour rather than automated processes.

3.1. Activity unification and temporal aggregation

To enable consistent analysis across the different types of contributions, we first grouped in uniform formatting all recorded contribution actions (commits, pull requests, issues, and comments) into a common data structure. This was necessary because each activity type is recorded differently on GitHub and conveys distinct contextual information. By consolidating these heterogeneous sources into a unified format, where each record includes a contributor identifier, a timestamp, and an activity type, we were able to compute contributor-level metrics in a coherent manner over time. This consolidation facilitated two key dimensions of analysis: contributor-based aggregation (i.e., to examine how effort is distributed among participants) and time-based aggregation (i.e., to observe how these patterns evolve).

The final dataset covers a 54-month period and comprises 6741 human-generated activity records from 96 distinct contributors.

All data were aggregated on a monthly basis. This choice, rather than using weekly or yearly intervals, was informed by three considerations. First, monthly aggregation offers a balance between granularity and interpretability: it is fine-grained enough to capture medium-term variations in activity, yet coarse enough to avoid excessive noise from short spikes in participation. Second, monthly intervals align with typical reporting and review cycles in open-source governance contexts. Third, from a statistical perspective, monthly aggregation ensures a sufficient number of observations per contributor, supporting stable estimates of activity distribution and inequality metrics. Following the unification of activity types, we computed both raw activity counts and a composite weighted activity score for each contributor on a monthly basis. The weighted score was introduced to reflect the varying levels of technical complexity and review effort associated with different

[2] https://www.cryptomiso.com

types of contributions. We assigned weights to each activity type based on qualitative judgements informed by prior research on software engineering practices and collaborative development [32]:

- **Pull requests (PRs):** weight = 5. Pull requests typically involve substantial code changes, peer review, and merge decisions. They are more structured and subject to formal approval processes.
- **Commits:** weight = 3. Commits represent direct modifications to the codebase but may vary in complexity and are not always reviewed prior to integration.
- **Issues:** weight = 2. Issues include the identification of bugs, suggestions for enhancement, or user feedback. They initiate problem-solving processes, but do not constitute implementation work.
- **Comments:** weight = 1. Comments support discussion and coordination, often in response to other activities. While essential for collaboration, they generally require the least technical input individually.

These weights were applied to each activity type to calculate a monthly *weighted activity score* for each contributor, alongside the raw counts. Including both measures, raw and weighted, enabled us to distinguish between the quantity of engagement and the depth of involvement. This distinction is particularly relevant when assessing decentralisation: a contributor who frequently comments, but does not engage in implementation, should not be considered equivalent to one who submits pull requests that are subject to review and integration. While the weighting scheme is necessarily heuristic, it seeks to approximate the typical effort and influence associated with each type of contribution within the collaborative development process.

3.2. Contributor distribution metrics

To evaluate the extent to which development activity is decentralised across contributors, we applied three widely used distributional metrics: the Gini coefficient [25], the Herfindahl–Hirschman Index (HHI) [26], and the Contributor Share Concentration (Top 5 and Top 10). Each metric was selected to capture a different facet of distributional imbalance, drawing on established applications in economics, social stratification studies, and software engineering research on open-source development [28–30].

The *Gini coefficient* measures inequality in the distribution of activity across contributors. A Gini value of 0 corresponds to perfect equality, where all contributors are equally active, whereas values closer to 1 indicate greater disparity, with a small number of individuals accounting for a substantial share of activity [25]. This measure is particularly suitable for detecting structural imbalances, even in projects with a large number of contributors, and enables comparisons across teams of varying sizes. We computed the Gini coefficient separately for each activity type (e.g., commits, PRs) as well as for total and weighted activity aggregates. For a given set of observations with values $x_1, x_2, \ldots, x_n$ in ascending order, the Gini coefficient is calculated as:

$$G = \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} |x_i - x_j|}{2n \sum_{i}^{n} x_i}\ , \tag{1}$$

where $\bar{x}$ is the mean of the distribution.

To complement the Gini coefficient, we also calculated the share of total contributions made by the top 5 and top 10 most active contributors. While the Gini index provides a summary of the overall distribution, it may obscure the specific influence of the most active individuals. The top-N contributor share offers a more direct perspective on whether a small subset of contributors is responsible for a substantial portion of activity. This is particularly relevant in assessing decentralisation: for instance, if a project is described as community-driven but a majority of contributions originate from a small group of contributors, the strength of that claim warrants scrutiny.

The final metric, the *Herfindahl–Hirschman Index (HHI)* [26], is derived from industrial organisation economics, where it is used to evaluate market concentration. In this context, the index is defined as the sum of squared contribution shares across all contributors:

$$\mathrm{HHI} = \sum_{i=1}^{n} \left(\frac{x_i}{\sum_{j=1}^{n} x_j} \right)^2 \times 10^4\ ,$$

where x_i is the activity count (or weighted score) for contributor i. HHI values range from close to zero (indicating a large number of equally active contributors) to 10,000 (indicating complete concentration). Although originally developed to assess firm dominance in market contexts, the HHI has been adapted in software engineering research to capture patterns of dominance in collaborative projects. Compared to the Gini coefficient, it places greater emphasis on larger individual shares and is more sensitive to changes among the most active contributors.

All three metrics were computed across the complete dataset, as well as separately for each calendar month. Monthly computation enabled us to examine how contributor concentration changed over time, allowing us to assess whether development became more or less centralised during significant events such as governance changes, security incidents, or protocol upgrades. The use of multiple metrics contributes to robustness: when findings are consistent across the Gini coefficient, HHI, and top-N contributor shares, they are less likely to reflect artefacts of any one metric and more likely to capture underlying structural characteristics of the contributor base.

3.3. Sensitivity analysis of activity weights

To address potential concerns regarding the selection of weights in our weighted activity score calculation, we conducted a sensitivity analysis to assess the robustness of our findings.

Predefined scenario analysis. We tested six theoretically motivated weighting scenarios to evaluate the robustness of our findings. Table 1 presents the specific weights assigned to each activity type across these scenarios. Our original weights (PR=5, Commits=3, Issues=2, Comments=1) balance code contributions with collaborative activities. The code-centric scheme increases weights for commits (6) and pull requests (8) while minimising collaborative activities (Issues=1, Comments=0.5). Conversely, the collaboration-centric approach emphasises discussion and coordination by increasing weights for issues (3) and comments (2) while reducing code-related weights. We also tested equal weights (all activities=1), a pull request-dominant scheme (PR=10), and an issue-focused configuration (Issues=5). Across all scenarios, the Gini coefficient ranged from 0.797 to 0.861, consistently indicating a high level of contribution concentration. Notably, the collaboration-centric configuration yielded a higher Gini coefficient (0.856) than our original scheme (0.823), while the code-centric approach produced the lowest concentration (0.797). This suggests that the overall pattern of concentration persists regardless of whether code contributions or collaborative activities are prioritised, though emphasising collaborative activities actually increases measured inequality.

Monte Carlo simulation. We performed 10,000 iterations with weights drawn from a uniform distribution U(0.5, 10), while maintaining the logical constraint that pull requests should not receive a lower weight than comments. This exploration of the parameter space revealed consistent results. The Gini coefficient exhibited a mean of 0.842 with a standard deviation of 0.018, yielding a 95% confidence interval of [0.799, 0.865]. Similarly, the top 1% of contributors accounted for a mean share of 27.88%, with a 95% confidence interval of [18.54%, 32.18%]. These narrow confidence intervals indicate that the main concentration patterns are not artefacts of specific weight assignments.

Table 1
Weight configurations used in sensitivity analysis. The Gini column shows the Gini coefficient of weighted activity distribution. The Top 1% column indicates the share of total weighted activity contributed by the top 1% of contributors, ranked by weighted activity (with 96 contributors, this corresponds to the single most active contributor).

Scenario	Pull Requests	Commits	Issues	Comments	Gini	Top 1%
Original	5	3	2	1	0.823	24.06%
Code-centric	8	6	1	0.5	0.797	17.10%
Collaboration-centric	2.5	2	3	2	0.856	30.69%
Equal weights	1	1	1	1	0.861	31.72%
PR-dominant	10	2	1	0.5	0.800	16.40%
Issue-focused	3	2	5	2	0.850	29.38%

Rank correlation analysis. We computed Spearman rank correlations between contributor rankings generated under different weighting schemes. Correlation values ranged from 0.877 to 0.994, with most exceeding 0.9. This indicates that the identification of core contributors remains stable across weighting configurations. Whether emphasis is placed on implementation or collaboration, the same contributors consistently appear among the most active participants in the project.

Fig. 1 presents the results of the sensitivity analysis based on 10,000 Monte Carlo simulations using randomly generated activity weights. The top panels illustrate the distribution of the primary concentration metrics. Panel (a) shows a right-skewed distribution of the Gini coefficient, with a mean of 0.842 and a standard deviation of 0.018. Panel (b) displays the distribution of the top 1% contributor share, which is also right-skewed and exhibits a broader spread, with a mean of 27.88%. In both cases, the red dashed lines represent the values obtained using the original weighting scheme (Gini = 0.823, Top 1% = 24.06%), which fall below their respective means, hence indicating that the baseline estimates are conservative.

Panel (c) compares the distributions of four concentration metrics using box plots. The Gini coefficient exhibits the least variability, whereas the contributor-share metrics (1%, 5%, and 10%) display increasing sensitivity to changes in weight configurations. Despite this variation, all metrics suggest a consistently uneven distribution of activity across simulations: the median Gini coefficient exceeds 0.84, and the median top 1% contribution remains above 27%.

Panel (d) illustrates the joint distribution of the Gini coefficient and the top 1% contributor share across all simulations. A strong positive—albeit slightly non-linear—association emerges between the two metrics. The red star indicates the outcome using the original weights, situated in the lower-left area of the distribution. This correlation pattern implies that higher values of one measure tend to be associated with higher values of the other, reinforcing the robustness of the observed concentration patterns. The placement of the original weighting scheme within the broader distribution further supports the interpretation that it provides a conservative estimate.

The consistency of results across a wide range of weighting schemes supports the validity of our methodological approach. Even in contrasting configurations — such as equal weights that treat all activity types identically, or code-centric weights that downplay collaborative actions — the overall pattern of contributor concentration remains evident. In the most evenly distributed scenario, the top 1% of contributors account for at least 18.54% of weighted activity, while in the most concentrated configuration, this share rises to 32.18%. Both figures indicate a notable imbalance, well above what might be expected in a fully decentralised development setting.

The original weighting scheme, which seeks to balance code-related and collaborative contributions, offers a representative measure of this distribution without disproportionately accounting for any specific activity type.

3.4. Interaction network construction

While contribution metrics such as commits and PRs capture direct technical activity, they do not account for the social dimension of development. To complement these measures, we constructed a contributor interaction network based on GitHub discussions. This network approximates patterns of collaboration and communication by modelling interactions between contributors through their shared participation in issue and PR threads.

In this network, each node represents a unique contributor. An undirected, weighted edge is added between two contributors if both have commented on the same issue or pull request, with the edge weight representing the number of issues/PRs on which they have co-participated. This connection does not imply direct dialogue, but rather co-participation in a shared context of problem-solving or design discussion. The underlying assumption is that contributors who engage in the same threads are likely to be aware of each other's work, with stronger weights indicating more frequent collaboration, thereby supporting informal coordination and knowledge diffusion within the project. This weighted approach is consistent with prior empirical software engineering research, where comment-based proximity and frequency of interaction have been shown to reflect meaningful social ties, particularly in distributed teams.

To analyse the resulting graph, we applied standard graph-theoretical centrality metrics:

- *Degree centrality* measures the number of distinct collaborators (i.e., other contributors to whom a node is connected). High degree centrality indicates broad involvement in conversations and may suggest a coordinating role.
- *Betweenness centrality* quantifies the extent to which a contributor lies on the shortest paths between other nodes. A high score indicates a bridging role between otherwise disconnected subgroups, which may be important for cross-module coordination or agreement-building.
- *Eigenvector centrality* captures influence beyond direct connections by assigning higher scores to contributors who are connected to other highly connected individuals. This metric is useful for identifying participants whose impact extends through indirect relationships.

Each of these centrality metrics offers a distinct perspective on collaboration: degree centrality reflects breadth of interaction, betweenness highlights strategic positioning, and eigenvector centrality captures embedded structural influence within the network [36,37].

This network-based perspective complements the distributional metrics by providing insight into how contributors engage with one another, rather than focusing solely on the volume of their contributions. It is particularly relevant when assessing decentralisation, as a development process characterised by a single, tightly connected cluster may reflect a different form of centralisation compared to one where contributions are dispersed but unevenly distributed. To further characterise the network structure, we applied core–periphery analysis to partition contributors into two groups: a densely connected core and a sparsely connected periphery. Contributors were classified based on their degree centrality, using the median centrality value as the threshold for core membership. We computed connection densities within the core, within the periphery, and between the two groups, as well as a core–periphery fitness score (defined as the difference between core density and periphery density) to measure the strength of this structural division.

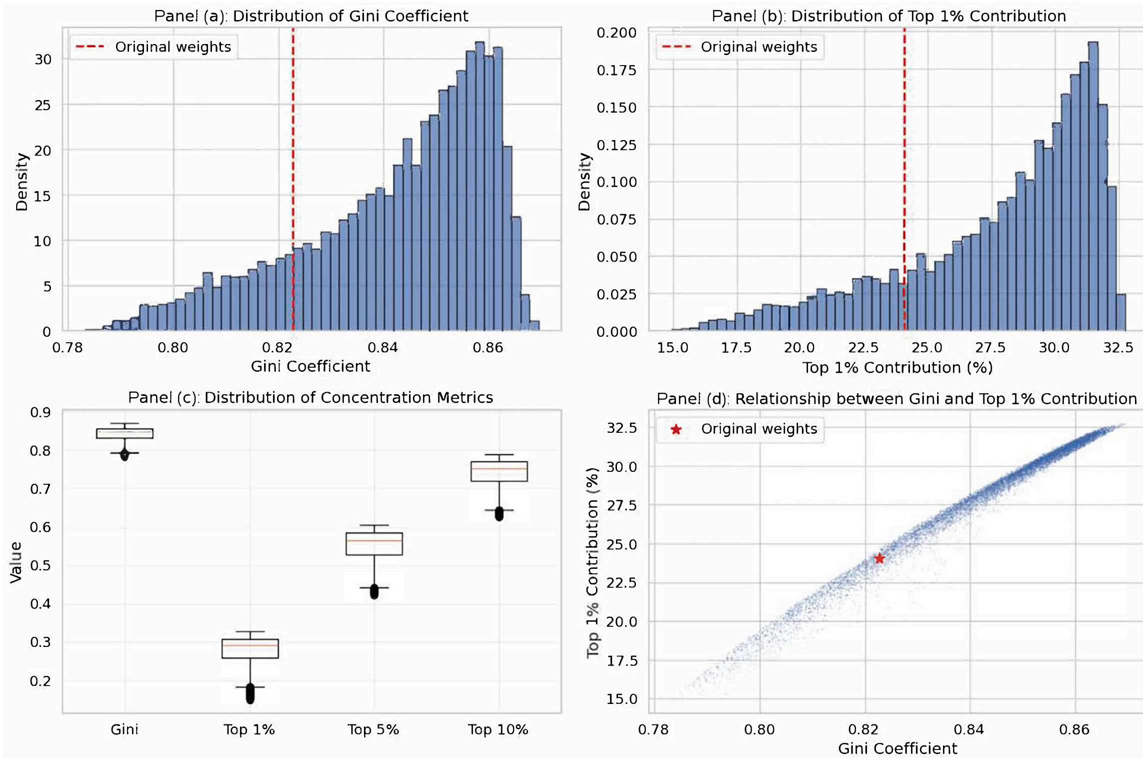

Fig. 1. Sensitivity analysis of activity weights. Top panels show the distribution of (a) Gini coefficients and (b) share of activity from the top 1% of contributors across 10,000 Monte Carlo iterations with random weights. (c) panel displays box plots for all concentration metrics. (d) panel illustrates the relationship between the Gini coefficient and top 1% contributor share, with the red star indicating results from our original weighting scheme.

3.5. Time-series analysis of development activity

To examine the relationship between development activity and protocol value, we employed time-series analysis techniques using Total Value Locked (TVL) as a proxy for protocol adoption and trust. TVL data was obtained from DefiLlama and aggregated to monthly averages to align with our activity metrics.

For this analysis, we focus specifically on commit activity rather than the weighted activity index used in concentration analysis. This distinction is deliberate: while the weighted activity index captures the breadth of contributor engagement (including discussion and coordination through comments), commit activity represents deployed code changes that directly affect protocol functionality. TVL, as a measure of capital locked in the protocol, responds primarily to technical improvements, security updates, and feature deployments rather than to discussion volume. Empirically, we observed that weighted activity (which is dominated by comments) shows near-zero correlation with TVL (Pearson $r = -0.08$), while commit activity demonstrates moderate positive correlation.

We applied three complementary analytical approaches:

Correlation analysis. We computed Pearson and Spearman correlation coefficients between TVL and monthly commit counts. The Spearman rank correlation provides robustness to outliers and non-linear relationships common in financial time series.

Cointegration testing. To assess whether TVL and development activity share a long-term equilibrium relationship, we applied the Engle–Granger two-step cointegration test [38]. This test determines whether two non-stationary time series move together over time, even if they diverge in the short term. We tested both raw TVL and log-transformed TVL, as financial series often exhibit exponential growth patterns. A statistically significant cointegration relationship ($p < 0.05$) indicates that the series share a common stochastic trend.

Lag correlation analysis. To investigate temporal precedence, we computed lagged cross-correlations between TVL and commit activity for lags ranging from −3 to +3 months. Negative lags indicate that commit activity leads TVL changes, while positive lags indicate that TVL leads commit activity. This analysis helps identify whether development efforts precede value accrual or vice versa.

Event analysis. We overlaid development activity patterns with major external events to examine how the protocol responds to ecosystem shocks and governance decisions. Events were categorised into two types: (1) market events including the Terra/Luna collapse (May 2022), Ethereum Proof-of-Stake merge (September 2022), FTX collapse (November 2022), and Eigenlayer mainnet launch (April 2024); and (2) governance decisions including major DAO votes and protocol upgrades. This qualitative analysis complements the quantitative measures by contextualising activity patterns within the broader DeFi ecosystem.

4. Results

In Fig. 2 we show the overall Lido monthly activity. The number of unique contributors per month is relatively low, with a long-term average of 7.2. Even during peak months, the number of unique contributors rarely exceeds 20, suggesting that development is concentrated within a small group. Certain periods, such as late 2021 and mid-2024, show minimal activity and reduced contributor engagement, highlighting the uneven distribution of participation over time.

These findings point to a pattern of centralised contribution, in which the majority of activity, particularly high-weighted actions such as pull requests, is carried out by a limited subset of contributors. When broader participation does occur, it tends to be associated with discussion-based actions rather than direct code implementation.

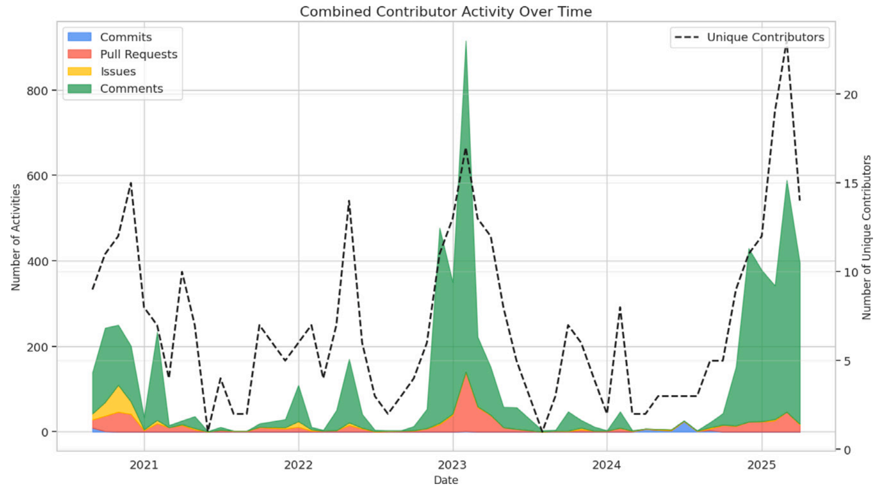

Fig. 2. Monthly development activity in Lido by type (stacked area), with unique contributors per month (dashed line, right axis). Comments dominate the total volume of activity. Spikes in contribution frequency often align with governance events or protocol upgrades.

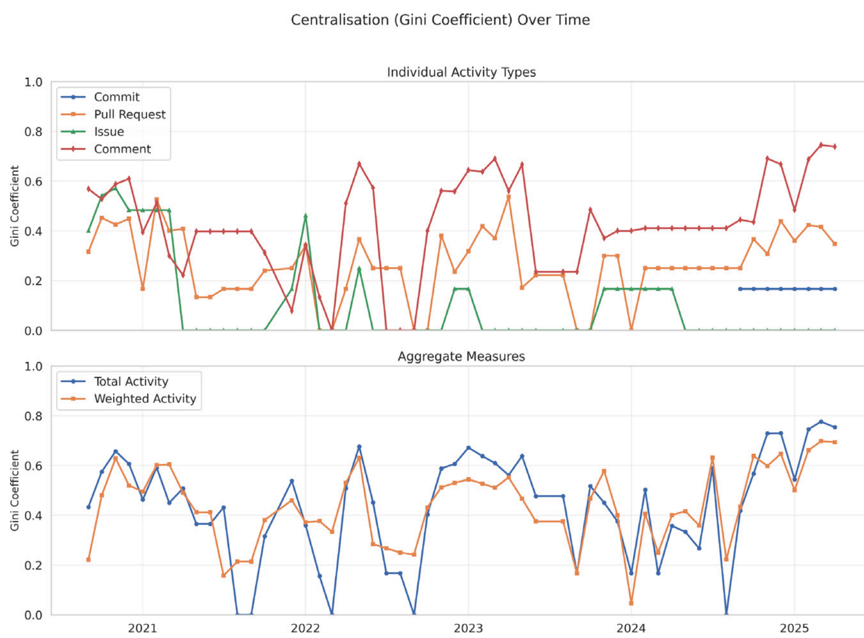

Fig. 3. Monthly Gini coefficients for each activity type in Lido. Higher values indicate greater inequality in the distribution of contributions.

4.1. Concentration of contributions

To quantify the observed concentration of contributions, we computed the Gini coefficient over time for each activity type, as well as for the total and weighted activity aggregates. Fig. 3 shows the monthly evolution of these metrics. The top panel displays Gini coefficients for individual activity types (commits, pull requests, issues, and comments), while the bottom panel presents aggregate measures (total activity and weighted activity). This separation reveals that while individual activity types exhibit considerable volatility, the aggregate measures show a more stable pattern with an upward trend in centralisation over time.

The Gini coefficient evolution provides empirical support for the “decentralisation illusion” hypothesis within DeFi development. Throughout the 54-month observation period, centralisation metrics remain consistently high, with Gini coefficients typically ranging between 0.4 and 0.8 across different activity types, significantly above the 0.5 threshold commonly associated with moderate-to-high inequality. The temporal patterns lead to several observations. First, the higher volatility seen during Lido’s early development phase (2021) reflects the structural instability of contribution dynamics in the initial stages of DeFi projects, when governance mechanisms and contributor incentives may not yet be well established. Second, the recurring drops to zero in issue-related activity (green line) reflect months with either no issues created or issue activity limited to single contributors, highlighting the irregular nature of participation in technical discussions. Third, the gradual stabilisation of commit-related Gini values around 0.17 by 2025 suggests a modest trend towards broader distribution of code contributions, though still distant from the decentralised ideal often presented in DeFi governance narratives.

The sustained concentration in comment activity and total activity (in the 0.4 to 0.7 range) indicates that, despite decentralised governance through the Lido DAO, actual development dialogue and coordination remain concentrated among a small group of contributors. These temporal findings are consistent with the results of our core–periphery network analysis, and suggest that governance decentralisation does not necessarily imply decentralised development practices.

Table 2
Phase boundary sensitivity analysis. All configurations show monotonic increase in weighted activity Gini coefficient from earlier to later phases.

Configuration	Phase 1	Phase 2	Phase 3	Phase 4
Original (Jun 2021, Jan 2023)	0.686	0.713	0.817	–
Shifted +3 months (Sep 2021, Apr 2023)	0.696	0.743	0.829	–
Shifted –3 months (Mar 2021, Oct 2022)	0.603	0.686	0.817	–
Two phases (Jul 2022)	0.706	0.826	–	–
Four phases (yearly)	0.686	0.689	0.690	0.831

Throughout the 54-month period, the Gini coefficients for total and weighted activity generally remain above 0.5, indicating a consistent degree of inequality in the distribution of development tasks. In some months, values approach or exceed 0.7, suggesting that a small number of contributors is responsible for the majority of activity. This pattern is particularly evident in the comment and weighted activity categories, where inequality increases during periods of heightened overall engagement. Commits and PRs also show periods of concentration, though with greater variance. This is likely due to their lower frequency and the smaller number of contributors involved. Issues tend to exhibit lower Gini values when present, although data are missing or sparse in several months, resulting in discontinuities in the time series. Taken together, the Gini time series supports the earlier descriptive findings: development activity in Lido appears to be concentrated among a limited group of contributors, with fluctuations in active participation having only a modest effect on the overall distribution of effort.

To consider whether temporal aggregation might obscure structural change in contributor concentration, we conducted an additional analysis dividing the 54-month period into separate lifecycle phases. This allows us to assess whether the high overall Gini coefficient of 0.82 reflects temporary concentration during early development, or whether centralisation remains throughout the protocol's evolution. We divided the 54-month observation period into three approximately equal phases: *Bootstrap* (September 2020–May 2021, 9 months), *Growth* (June 2021–December 2022, 19 months), and *Maturity* (January 2023–February 2025, 26 months). While these boundaries do not correspond directly to specific protocol milestones, they provide a framework for assessing whether centralisation patterns change as the protocol evolves. The unequal phase durations reflect the natural progression of the project, with a shorter initial development period followed by extended growth and operational phases. For each phase, we computed separate concentration metrics, including Gini coefficients, unique contributor counts, and top-N contributor shares.

The distribution of activity among top contributors follows a similar pattern. The share of weighted activity attributable to the most active contributor rose from 16.9% in the bootstrap phase to 31.4% during maturity. Meanwhile, the average number of monthly active contributors remained low throughout (9.2, 5.2, and 7.9 in each phase respectively), suggesting that the growth in the contributor base did not result in sustained engagement.

These findings suggest that Lido's development became more concentrated over time, rather than shifting towards broader participation. The higher concentration observed in the maturity phase, relative to the bootstrap phase, points to a structural trend rather than a temporary feature of early development.

To validate that this increasing centralisation pattern is not an artifact of our specific phase definitions, we tested four alternative temporal partitions: boundaries shifted forward by three months, boundaries shifted backward by three months, a two-phase division at the observation period midpoint, and a four-phase yearly division. Table 2 presents the results. Across all five configurations (including the original), the weighted activity Gini coefficient increased monotonically from earlier to later phases. This consistency confirms that the observed centralisation trend is robust to boundary selection.

This raises questions about whether, without deliberate structural changes, DeFi protocols may tend towards greater centralisation as they evolve, complicating claims about the relationship between protocol growth and development decentralisation.

To further examine the temporal dynamics of contributor activity, we tracked the weighted monthly contributions of individual contributors over time. Fig. 4 presents a mountain graph visualisation that illustrates how development effort is distributed across contributors throughout the observation period. The figure supports the concentration patterns observed in the Gini analysis. During a period of heightened development in early 2023, Contributor A recorded a peak of over 300 weighted contributions in a single month. This peak coincides with key phases of protocol development, suggesting that substantial portions of technical work were undertaken by a single contributor.

The visualisation highlights the top ten individual contributors (A through J), while the yellow "Others" category, representing 86 additional contributors, shows a consistent pattern: although numerically dominant, these contributors account for only intermittent and limited activity. The figure captures not only static inequality, but also dynamic shifts in contribution patterns. Development activity tends to cluster around a small group of core contributors (particularly A, B, C, and D), while most participants remain on the periphery in terms of sustained engagement. Periods of high activity, especially those corresponding to major protocol updates, appear to be primarily driven by individual contributors rather than broad-based collaboration.

To illustrate the distribution of development effort, we plot the Lorenz curve for weighted activity in Fig. 5. The Lorenz curve displays the cumulative percentage of contributions (on the y-axis) against the cumulative percentage of contributors (on the x-axis), with contributors ordered by their contribution volume in ascending order. In a scenario of perfect equality, the distribution would follow the diagonal line $y = x$ (the line of equality). In practice, empirical distributions typically fall below this line, with the area between the curve and the diagonal representing the extent of inequality.

The curvature of the Lorenz curve indicates a markedly unequal distribution: a small proportion of contributors are responsible for the majority of substantive development activity. The corresponding Gini coefficient is 0.82, consistent with the values observed in the temporal Gini plot, and reinforces the finding that a limited group of individuals carries out most of the core implementation work. This static representation complements the monthly metrics and supports the observation that decentralised governance structures do not necessarily translate into decentralised development activity.

RQ1 asked: How is development activity distributed among contributors in Lido? *Answer: Our analysis reveals that development activity is highly concentrated among a small subset of contributors. The Gini coefficient for weighted activity is 0.82, indicating inequality in the contribution distribution. Only 96 unique human contributors participated over 54 months, with an average of just 7.2 active contributors per month. The core group accounts for 98.1% of all development activities. This concentration persists across all activity types, with pull requests and commits showing particularly high centralisation.*

4.2. Contributor interaction network

To supplement the activity-based indicators, we constructed a contributor interaction network based on comments and examined its

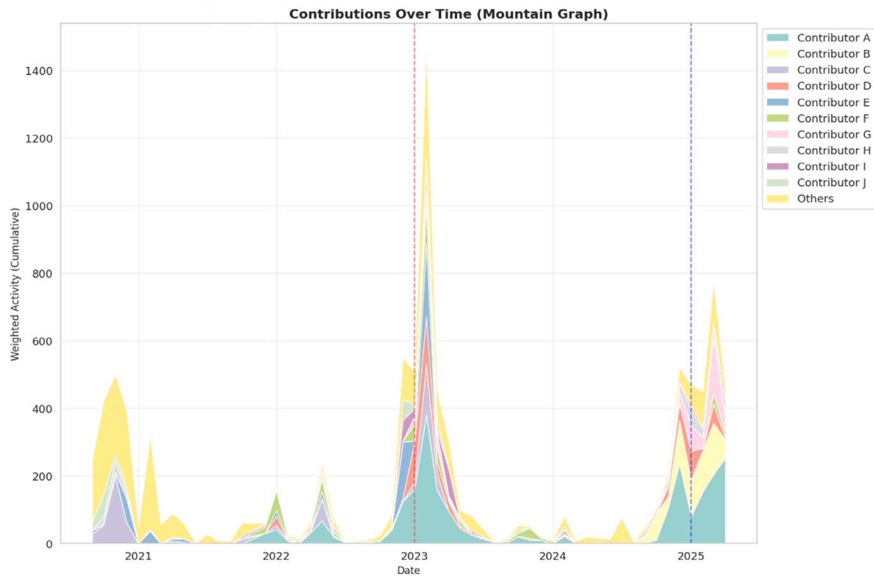

Fig. 4. Mountain graph visualisation of anonymised contributor contributions over time. Contributor A represents the most active contributor (corresponding to the dominant peak in early 2023), while the "Others" category contains the remaining 86 contributors who collectively contribute minimal development effort. The visualisation demonstrates extreme temporal centralisation, where a single contributor (Contributor A) dominated development activity during critical protocol development phases.

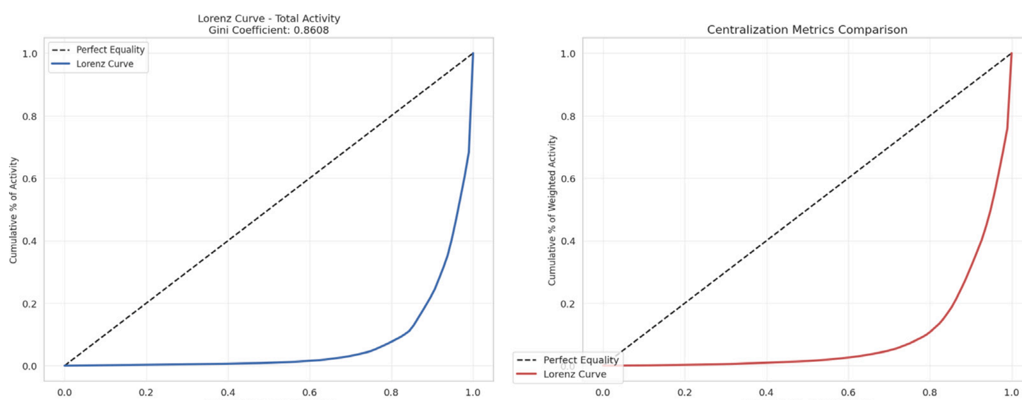

Fig. 5. Lorenz curve of total (left) and weighted (right) contributor activity in Lido. The high curvature reflects a strong concentration of contributions.

structural characteristics. Fig. 6 presents key metrics across four dimensions: centralisation, community structure, cohesion, and network distance.

The network comprises 96 contributors and 203 undirected, weighted edges, where each edge represents co-participation in issue or pull request discussions and edge weights indicate the number of shared discussions.

The network is fragmented into 16 disconnected components, with the largest component containing 81 contributors (84.4% of the network).

Centrality metrics indicate high concentration: the Gini coefficient computed over node degrees is 0.576, while betweenness centrality yields a higher Gini coefficient of 0.854, suggesting that a small number of contributors mediate a substantial proportion of information flow. Notably, the eigenvector centrality Gini of 0.551 also indicates concentration, contradicting expectations of distributed influence. The degree centralisation score of 0.397 further confirms the presence of a hierarchical structure.

Regarding community structure, analysis of the largest component identified 7 distinct sub-groups with a modularity of 0.182, indicating weak community separation. The average community size Gini of 0.406 suggests slightly uneven distribution of members across communities.

Cohesion metrics reveal a sparse network. The network density of 0.0445 (4.5%) indicates limited overall connectivity. While the clustering coefficient of 0.353 suggests some local collaboration, the fragmentation into 16 components undermines network-wide cohesion. Within the largest component, a diameter of 6 and average shortest path of 2.76 indicate that connected contributors can reach each other, but only through key intermediaries. These findings reveal that development coordination operates through isolated clusters with limited cross-cluster collaboration. The combination of high betweenness centralisation (0.854) and network fragmentation demonstrates that Lido's development depends on a few key coordinators who bridge otherwise disconnected groups, reinforcing the centralisation evident in contribution metrics.

4.2.1. Community structure visualisation

We visualised the collaboration network of the largest connected component (81 contributors, 84% of the network) with nodes coloured by their detected community, as shown in Fig. 7. The Louvain method [39] identified seven communities within this component. Community 0 (blue) is the largest, encompassing 30 contributors (37% of the component). This is followed by Community 1 (orange) with 16 contributors, Community 2 (red) with 14 contributors, and progressively

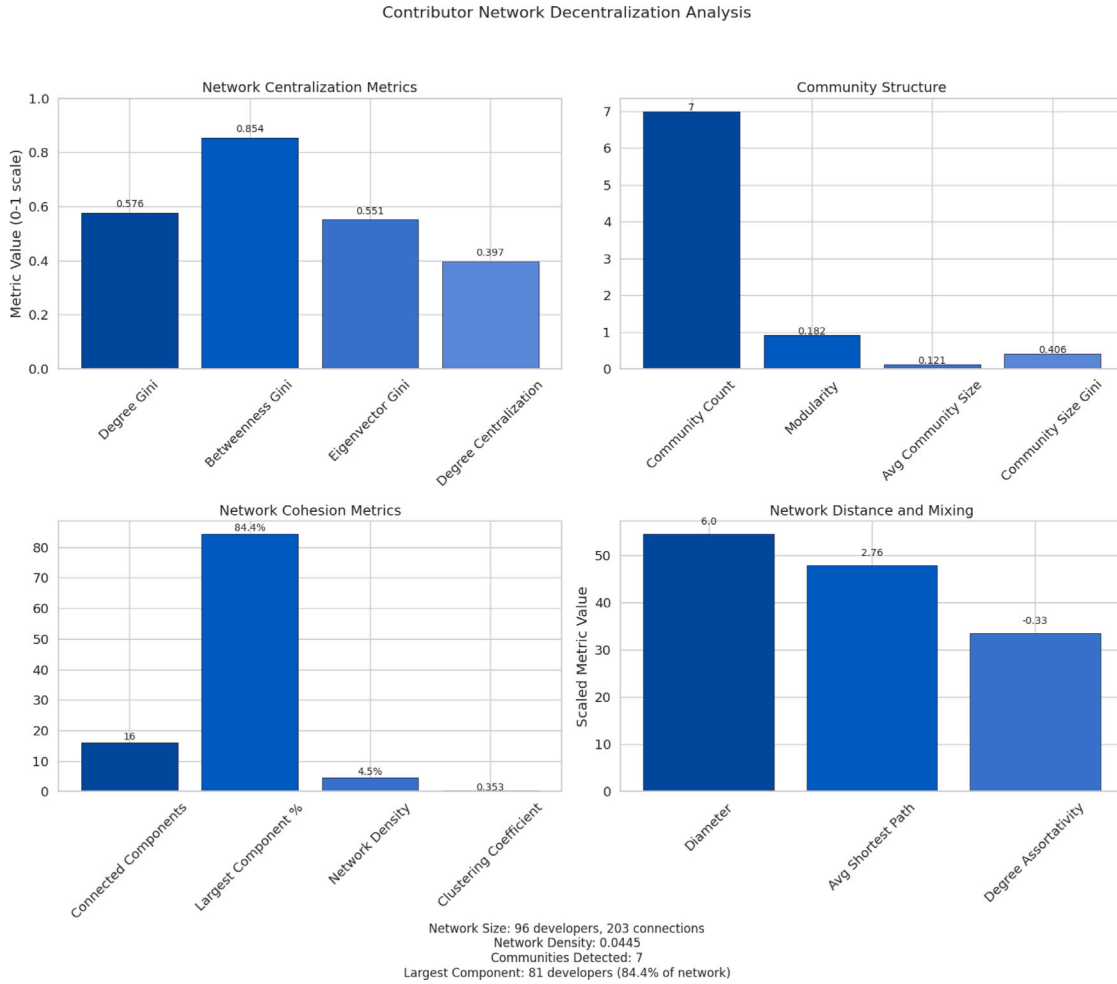

Fig. 6. Summary of collaboration network metrics for Lido.

smaller communities: Community 3 (8 contributors), Community 4 (7 contributors), Community 5 (4 contributors), and Community 6 (2 contributors).

The visualisation reveals a clear hub-and-spoke topology, with development coordination flowing through a small set of highly connected contributors. Contributors O1 and F1 occupy the most central positions, serving as primary bridges between communities. Contributors A, F, E, C, and O also function as important hubs, each connecting multiple peripheral contributors who would otherwise be isolated. This hierarchical structure is reflected in the network's high betweenness centrality Gini coefficient of 0.854.

The modularity value of 0.182 indicates weak community separation, suggesting that despite the visual clustering, communities are not functionally independent. Instead, they depend on central coordinators for cross-community collaboration. Most contributors connect to only one or two others, relying on the hub contributors to access the broader network. This structure aligns with earlier findings of concentrated development activity, where collaborative interactions are mediated by a small number of highly connected contributors rather than through distributed peer-to-peer connections typical of decentralised systems.

4.2.2. Contributor network evolution

To examine changes in decentralisation over time, we analysed the evolution of the contributor collaboration network throughout the lifespan of the Lido project. The dataset was divided into six non-overlapping time intervals, and core network statistics were computed for each period.

Fig. 8 presents key quantitative trends. The top-left panel shows the number of active contributors and their connections over time. While the number of contributors fluctuates modestly (between 9 and 24), the number of connections varies dramatically (from 10 to 91), indicating phases of intensified collaboration. The top-right panel depicts the network's density, which peaks at 0.38 in the 2023-10 period before declining, suggesting a phase of tighter coordination among contributors.

The bottom-left panel displays the Gini coefficient of node degrees, which serves as a proxy for the centralisation of the contributor network. The data reveal significant fluctuations: a notable drop to 0.19 during 2021-06 (the most decentralised period) followed by increases reaching 0.47 by 2022-12. The bottom-right panel shows a critical finding: during the 2021-06 period, only 67% of contributors belonged to the largest connected component, indicating network fragmentation. This fragmentation coincides with the lowest activity period (9 contributors, 10 connections).

Fig. 9 presents the temporal evolution of Lido's contributor collaboration networks, segmented into six time windows. Each subgraph depicts the co-engagement network formed by contributors interacting through issues, pull requests, or comments.

In the early phase (2020-09 to 2021-06), the network shows 21 contributors with 51 connections and Gini coefficient of 0.413. The structure reveals several high-degree nodes acting as coordination hubs.

The 2021-06 to 2022-03 interval exhibits the most dramatic change: contributors drop to 9, connections to just 10, and the network fragments (only 67% in the largest component). Despite this contraction,

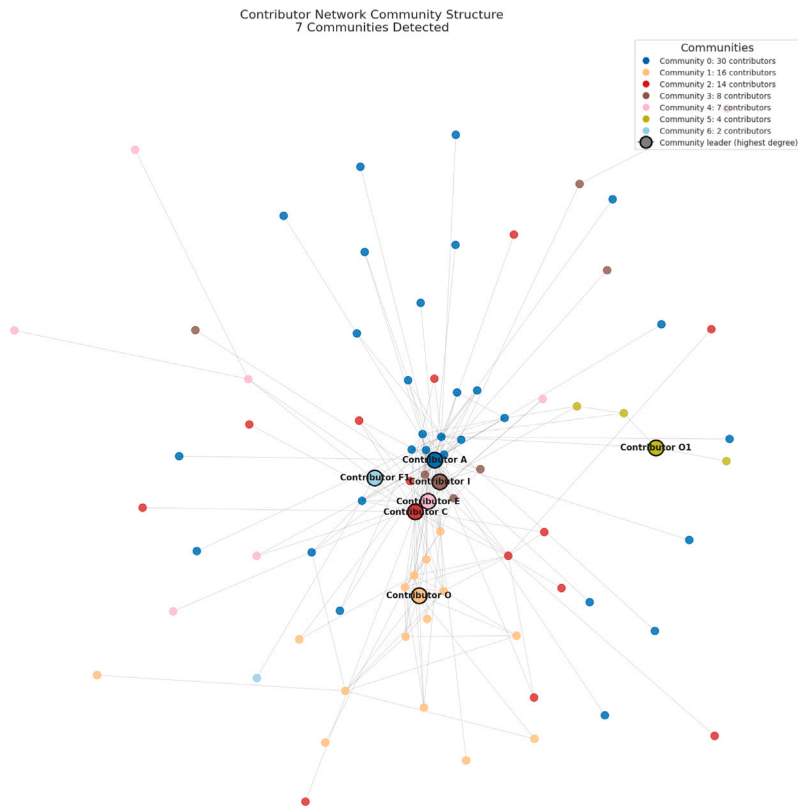

Fig. 7. Community structure of the largest connected component (81 contributors) in the Lido contributor network. Seven communities are identified, with a clear hub-and-spoke topology centred on Contributors O1, F1, A, F, E, C, and O. Node size represents degree centrality. The visualisation shows 84% of the network; the remaining 15 contributors exist in isolated components.

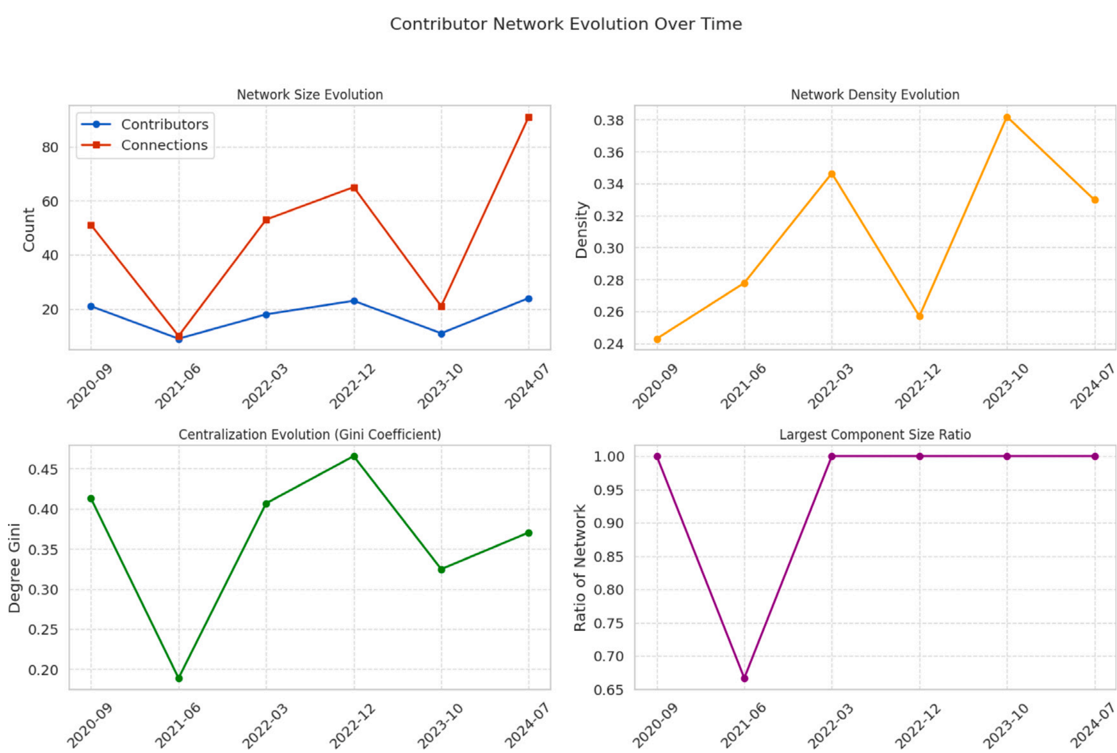

Fig. 8. Contributor Network Evolution Over Time. Top-left: count of contributors and connections. Top-right: evolution of network density. Bottom-left: degree-based Gini coefficient (centralisation). Bottom-right: share of contributors in the largest connected component.

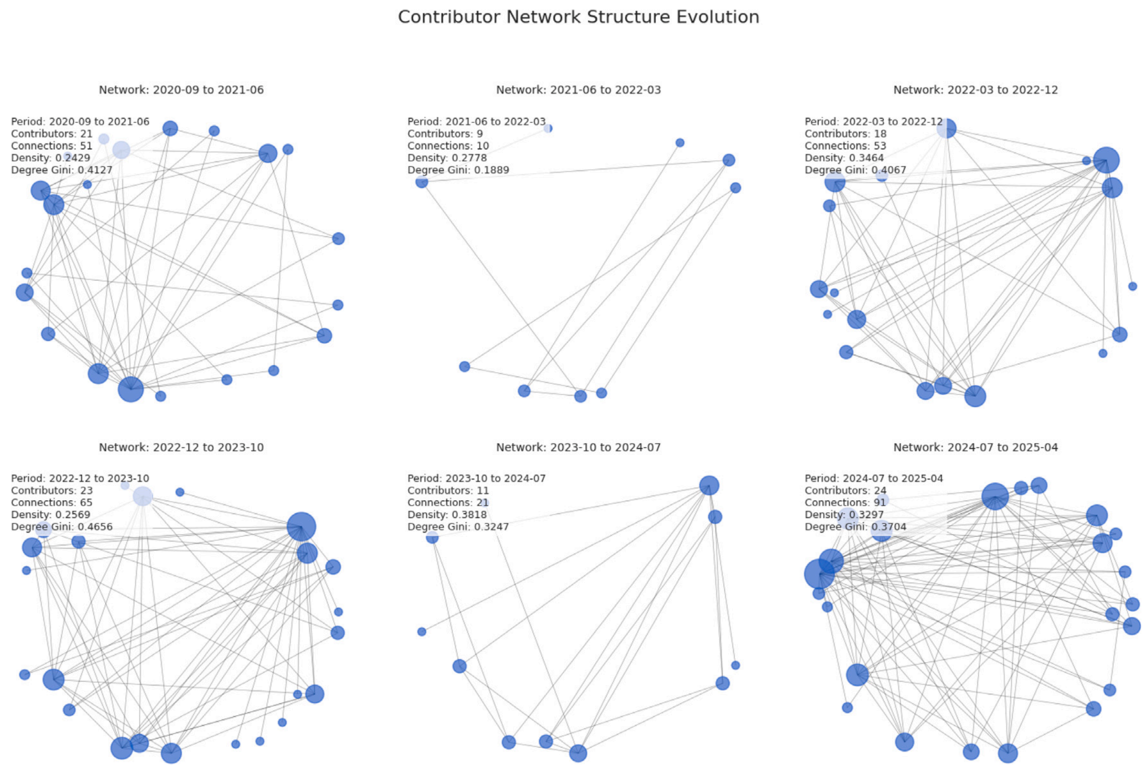

Fig. 9. Contributor Network Structure Evolution. Each subgraph represents the collaboration network within a time interval. Node size corresponds to degree centrality (i.e., number of links). Denser graphs indicate greater contributor interconnection. Degree Gini is shown for each period.

the Gini coefficient falls to 0.189—the lowest across all periods—suggesting more equal participation among the remaining active contributors.

The 2022-03 to 2022-12 period marks a recovery phase with 18 contributors and 53 connections. However, centralisation increases (Gini = 0.407), and the network reconnects to 100% cohesion. The visual structure shows renewed hub-and-spoke patterns.

The subsequent periods (2022-12 through 2025-04) show continued growth and fluctuation. The 2022-12 to 2023-10 interval reaches highest degree Gini (0.466) with 23 contributors. The network then slightly decentralises in 2023-10 to 2024-07 (Gini = 0.325) before rising again in the final period (Gini = 0.370) with the highest connectivity observed (24 contributors, 91 connections).

These temporal patterns reveal that Lido's development network experiences cycles of expansion and contraction, with centralisation generally increasing during growth phases. The 2021-06 fragmentation period likely reflects a transitional phase in the protocol's development, possibly corresponding to major architectural changes or governance restructuring.

Building on the temporal analysis of decentralisation, we now examine the structural composition of the contributor network using a core–periphery model [40]. Fig. 10 shows the result of this analysis.

The partition reveals an exact 50-50 split between core (48 contributors) and periphery (48 contributors). However, the distribution of development activity is highly asymmetric: 98.1% of all interactions originate from core contributors, while peripheral contributors account for only 1.9%. This extreme imbalance reinforces previous findings of concentrated participation and suggests that half of the network contributors have minimal engagement in development activities.

To address whether the high overall concentration is merely an artifact of including minimally active peripheral contributors, we computed the Gini coefficient within each group separately. The core group (48 contributors) exhibits a Gini coefficient of 0.687, while the periphery group shows 0.418. The high within-core Gini indicates that even among active contributors, development effort is unevenly distributed. This confirms that centralisation is not solely driven by the presence of inactive peripheral contributors, but reflects structural inequality within the productive core itself.

The connection density matrix reinforces this distinction. Core-to-core connections exhibit relatively high density (0.151), while core-to-periphery links are sparse (0.014) and periphery-to-periphery connections are essentially absent (0.000). This pattern indicates that peripheral contributors rarely interact with each other and have limited connections even to core contributors. The computed core–periphery fitness score (0.1507) confirms a clear structural separation, though not as extreme as the activity distribution might suggest.

This structure has critical implications for the protocol's sustainability. While the even split between core and periphery members might suggest balanced participation, the reality is that half the (peripheral) network contributes negligibly to development activity. The project's heavy reliance on 48 core contributors, with virtually no peer-to-peer collaboration among peripheral contributors, creates significant risk: the departure of even a small fraction of core contributors could severely impact development capacity.

These network characteristics—temporal fluctuations with increasing centralisation, fragmentation during transition periods, and stark core–periphery division—provide structural evidence for the concentration patterns observed in our distributional analysis. The architecture not only reflects existing centralisation, but likely reinforces it by possibly creating barriers to peripheral contributors becoming more engaged.

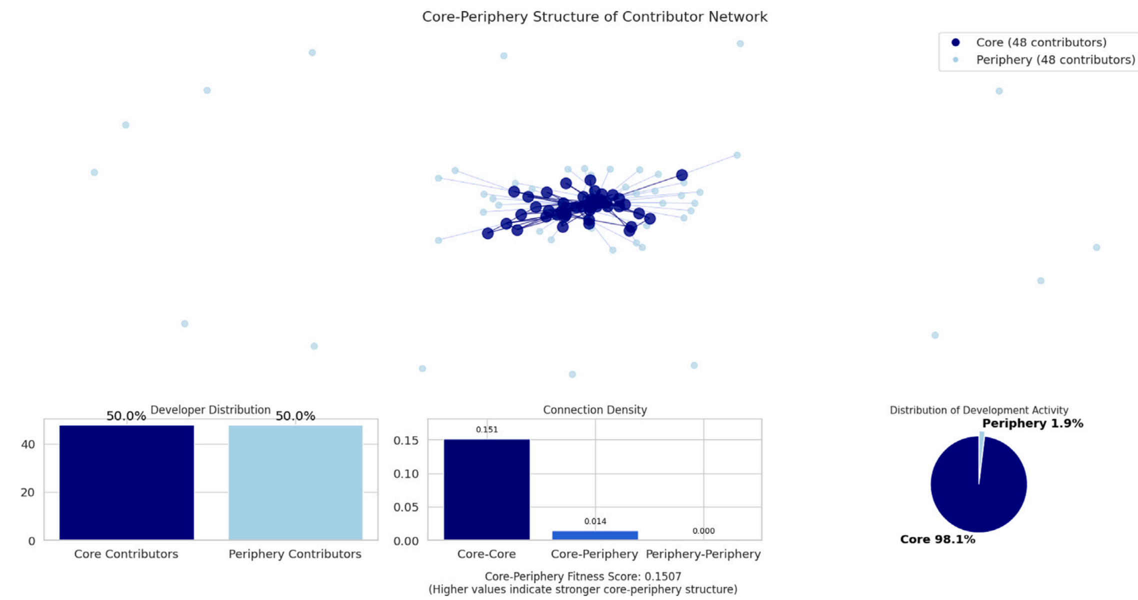

Fig. 10. Core–periphery structure of the Lido contributor network. Contributors are classified into core (dark blue) and periphery (light blue) groups. The bar chart shows an even contributor distribution (48 core, 48 periphery), yet a substantial imbalance in development effort: 98.1% of contributions stem from core contributors. Connection density is also skewed: most interactions occur within the core (density = 0.151), while cross-tier and periphery-periphery connections are minimal. The resulting core–periphery fitness score of 0.1507 confirms a moderate structural separation.

RQ2 asked: How do contribution patterns and collaboration structures evolve over time? *Answer: The temporal analysis demonstrates that development centralisation in Lido is not static but shows a general trend towards increased concentration over time. Monthly Gini coefficients fluctuate between 0.19 and 0.47, with the lowest centralisation occurring during a fragmentation period (2021-06) when the network contracted to just 9 contributors. The mountain graph visualisation reveals that development leadership shifts between a small set of core contributors (particularly Contributors A through O), while 86 peripheral contributors remain marginalised. Network evolution analysis shows cycles of expansion and contraction, with centralisation typically increasing during growth phases. The core-periphery structure reveals that while contributors are evenly split (48 core, 48 periphery), 98.1% of all development activity originates from the core group, demonstrating persistent concentration despite temporal fluctuations.*

4.3. External influence

To examine the relationship between Lido's protocol growth and its open-source development activity, we conducted correlation and cointegration analyses using monthly data on total value locked (TVL) and contributors' commit activity. As detailed in Section 3.5, we use commit activity rather than the weighted activity index for this analysis. While the weighted index captures overall engagement including discussion, commit activity represents code changes that directly affect protocol functionality. This distinction is empirically supported: weighted activity shows negligible correlation with TVL (Pearson $r = -0.08$, Spearman $\rho = -0.05$), while commit activity demonstrates moderate positive correlation (Pearson $r = 0.46$, Spearman $\rho = 0.55$). This suggests that protocol value responds to deployed code changes rather than to discussion volume.

The correlation analysis shows a moderate positive association between TVL and commit frequency, with a Pearson correlation coefficient of 0.46 and a Spearman rank correlation of 0.55. These values indicate that increases in commit activity tend to align with periods of higher TVL. However, the correlation is not strong enough to suggest a direct or exclusive dependency, implying the presence of additional influencing factors on both metrics.

To assess potential long-term relationships, we applied the Engle–Granger cointegration test. The results indicate that log-transformed TVL and commit activity are cointegrated at the 5% significance level (test statistic = -6.72, p-value < 0.001). This suggests a stable long-term equilibrium between protocol value and development intensity. While short-term deviations exist, the two variables tend to converge over time, reflecting an underlying structural link shaped by ecosystem expansion, protocol stability, and market dynamics.

Fig. 11 visualises these dynamics through two complementary analyses. The top panel displays normalised time series of TVL and commit activity from 2021 to 2025, revealing several periods of strong co-movement. Notable alignment occurs during the 2023 growth phase, where both metrics rise in tandem, while divergences appear during market stress periods such as mid-2022. The bottom panel presents the lag correlation analysis, examining correlation coefficients across a range of -3 to $+3$ months. The analysis reveals a clear asymmetric pattern: negative lags (where commits lead TVL) show stronger correlations than positive lags. Specifically, the peak correlation of 0.618 occurs at lag -3, substantially higher than the contemporaneous correlation of 0.46. This pattern strengthens progressively from lag 0 (0.46) to lag -1 (0.50), lag -2 (0.55), reaching maximum at lag -3 (0.618), before declining slightly at longer lags.

While short-term predictive causality is not supported — as confirmed by Granger causality tests — the moderate correlation and evidence of cointegration indicate that Lido's development activity and TVL are interconnected over the long term. The lag analysis provides crucial insight: the strongest correlation (0.618) at a 3-month lead time suggests that development efforts systematically precede value accrual in the protocol. This temporal relationship implies that sustained development activity may serve as a leading indicator for protocol growth, with technical improvements and feature implementations requiring approximately three months to translate into increased protocol adoption and value locked.

An examination of Lido's development activity, measured through commits and pull requests, also reveals meaningful patterns when aligned with major market disruptions and governance interventions. Beyond the expected increases in activity surrounding technical milestones, several external events appear to correspond with shifts in contributors' behaviour. The selected events include:

Market Events

- May 9–13, 2022: Terra/Luna collapse causing stETH depeg event

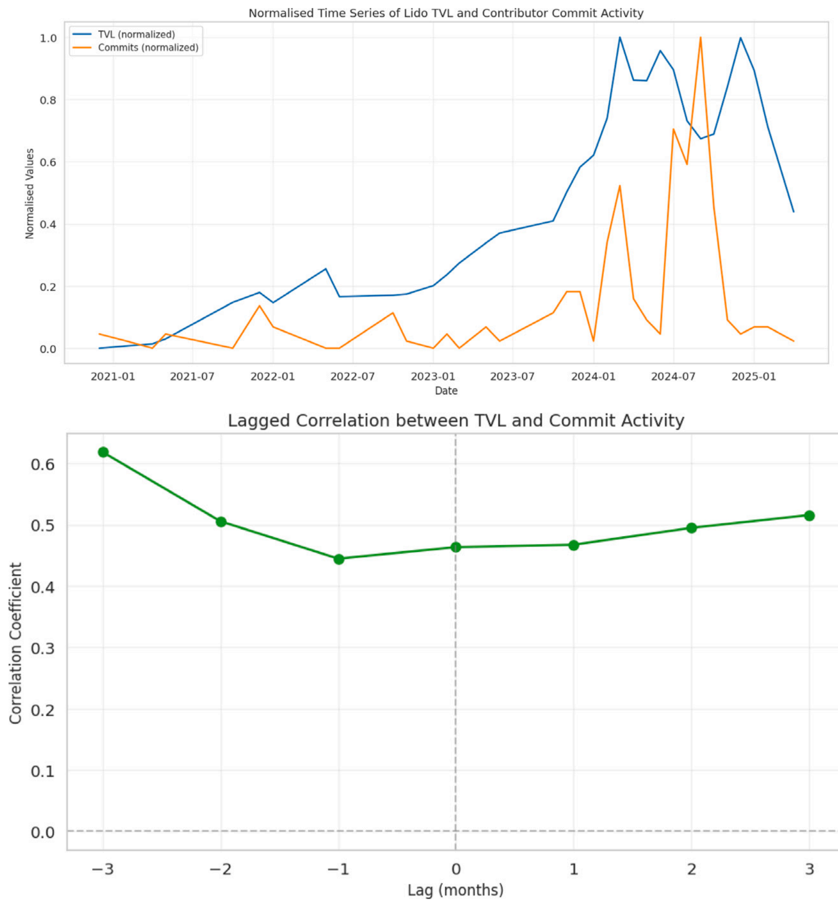

Fig. 11. Top: Normalised time series of Lido's Total Value Locked (TVL) and contributor commit activity from 2021 to 2025. The chart highlights periods of co-movement between protocol value and development efforts. Bottom: Lagged correlation between TVL and commit activity. The plot illustrates Pearson correlation coefficients for lags ranging from −3 to +3 months. The highest correlation (0.618) is observed at lag −3, indicating that commit activity leads TVL changes by approximately 3 months.

- September 15, 2022: Ethereum Proof-of-Stake merge
- November 8–11, 2022: FTX collapse affecting stETH liquidity
- March 28, 2023: Lido surpassed 30% of all staked ETH on Ethereum
- June 22, 2023: Exceeded $10 billion in TVL
- April 10, 2024: Eigenlayer mainnet launch enabling restaking
- May 23, 2024: SEC approval of spot Ethereum ETFs

Governance Decisions

- July 19, 2022: Dual Governance proposal passed
- December 21, 2022: Vote to phase out Lido on Solana
- February 7, 2023: Implementation of "Safe program" protocol insurance
- August 15, 2023: Governance vote on staking limits
- October 5, 2023: Passed LIP for Withdrawal Request System

During major market events, such as the Terra/Luna collapse in May 2022 and the FTX collapse in November 2022, we observe moderate increases in development activity, as illustrated in the bottom panel of Fig. 12.

The bottom panel specifically tracks market events against normalised development activity (commits and pull requests), revealing distinct spikes that align with crisis periods. For instance, the Terra/Luna collapse period (marked in May 2022) shows a notable surge in both commit and pull request activity, while the FTX collapse (November 2022) demonstrates a similar response pattern. These periods were followed by increased development activity focused on addressing emerging risks, maintaining protocol stability, or mitigating liquidity concerns. The visualisation clearly demonstrates that exogenous market shocks can lead to concentrated bursts of contribution as contributors work to sustain protocol functionality under stress.

Similarly, significant governance decisions correlate with elevated development activity, as demonstrated in the top panel of Fig. 12. The top panel presents governance-related events overlaid on the same normalised development metrics, allowing for direct comparison with market-driven responses. For example, the Dual Governance proposal passed in July 2022 corresponds with a pronounced spike in both commits and pull requests visible in the top panel, suggesting immediate implementation efforts. The vote to phase out Lido on Solana in December 2022 shows an even more dramatic increase in development activity, with the top panel revealing sustained elevated activity levels in the weeks following this decision. The "Safe program" implementation (February 2023) and the Withdrawal Request System LIP (October 2023) also show increased activity, though to varying degrees. These spikes coincide with DAO-initiated decisions, suggesting an association between governance activity and development intensity.

Overall, while routine development aligns with planned upgrades, both market volatility and governance actions are associated with periods of intensified activity. This dual dynamic — proactive execution of roadmap objectives and reactive responses to external disruptions — highlights the embeddedness of Lido's development process within broader ecosystem forces and DAO deliberations. These findings point to a model of protocol maintenance in which contributor activity is contingent on both internal planning cycles and governance and market events.

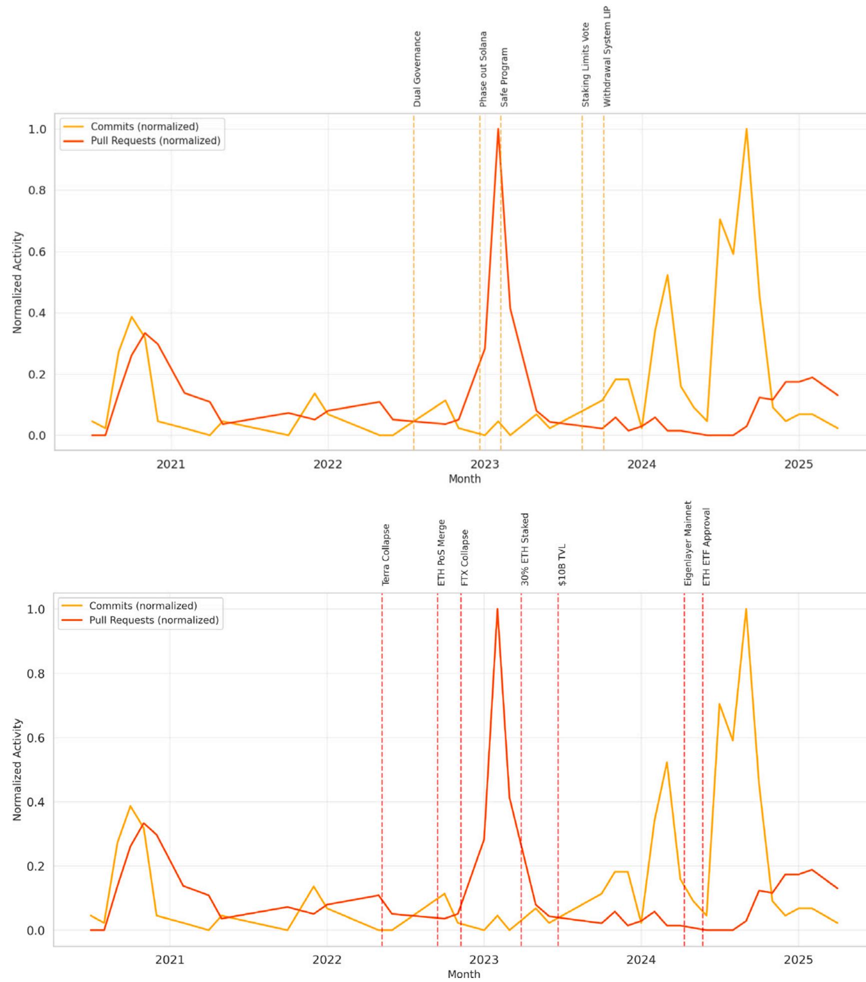

Fig. 12. Development activity patterns overlaid with key external events. Top: Governance decisions and their impact on development activity. Bottom: Market and ecosystem events including the Terra/Luna collapse, Ethereum Proof-of-Stake merge, FTX collapse, Eigenlayer mainnet launch, and Ethereum ETF approval. Vertical lines mark event dates; activity patterns show commits and pull requests normalised to [0,1] range.

> **RQ3 asked: How do external market shocks and governance decisions affect the distribution and temporal patterns of development activity in Lido?** *Answer: External factors demonstrate significant influence on development patterns, with TVL and commit activity showing moderate correlation (Pearson 0.46, Spearman 0.55) and long-term cointegration at the 5% significance level. The analysis reveals that commit activity leads TVL changes by 3 months, suggesting development efforts precede value accrual. Market crises such as the Terra/Luna collapse and FTX failure are followed by development spikes, while governance decisions including the Dual Governance proposal and Solana phase-out correlate with intensified development phases. These patterns indicate that while Lido's development responds to external stimuli, the responses consistently rely on the same concentrated group of core contributors, reinforcing rather than alleviating centralisation dynamics during critical periods.*

5. Discussion

The empirical evidence of development centralisation in Lido supports the 'decentralisation illusion' identified in prior DeFi research [17]. While the use of DAOs and token-based voting implies distributed control, our analysis shows that core development remains concentrated in a small group of contributors.

Our three research questions collectively reveal that not only is development highly concentrated RQ1, but this concentration intensifies as protocols mature RQ2, and external pressures that might be expected to democratise participation instead reinforce existing power structures RQ3. The 3-month lead time between development activity and TVL growth suggests that this concentrated contributor control is associated with protocol value changes, creating a feedback loop where those who control development also capture the benefits of growth.

The persistence of centralisation

The consistently high Gini coefficients observed over the 54-month period, reaching 0.82 for weighted activity, show that centralisation is a persistent structural feature, not a transitional phase. Moreover, this concentration has intensified over time, with the Gini coefficient rising from 0.686 during the bootstrap phase to 0.817 in the maturity phase. Despite Lido's size, maturity, and governance mechanisms, development leadership rotates among a narrow set of individuals rather than expanding to include peripheral contributors. This indicates that current incentive structures do not support broader participation.

The core–periphery network structure reinforces this view. Although core and peripheral contributors are numerically equal (48 each), their engagement is dramatically asymmetric: 98.1% of all development activity originates from the core group. The presence of moderate clustering (coefficient of 0.353) and negative degree assortativity (−0.331) indicates a disassortative network where high-degree core contributors connect primarily to low-degree peripheral contributors, reinforcing the hub-and-spoke topology that hinders upward mobility from periphery to core.

Governance-development disconnect

Governance decisions align with spikes in development activity, but these are consistently executed by the same small set of core contributors. In practice, decentralised governance acts as a coordination signal rather than a mechanism for redistributing development responsibility. The cointegration between Total Value Locked (TVL) and development activity further show this imbalance. Our analysis reveals that development activity leads TVL changes by approximately 3 months (correlation of 0.618 at lag −3), suggesting that while protocol value increases alongside development intensity, authority over that effort remains centralised. This feedback loop reinforces contributor concentration as the protocol matures.

Structural drivers and risks

Network analysis clarifies how centralisation is sustained. Contributors with high betweenness centrality dominate communication flows (betweenness Gini of 0.854), effectively acting as gatekeepers of technical knowledge and coordination. The low modularity (0.182) and the fragmentation into 16 disconnected components, with 84.4% of contributors in a single large component, limit the emergence of specialised or autonomous working groups.

These structures present operational risks. A small number of contributors oversee critical functions, creating a single point of failure. While the responses to the Terra/Luna and FTX collapses show that the core team can respond effectively, they also highlight the fragility of relying on a few individuals during periods of stress or parallel crises.

Conceptual and methodological implications

This study shows that decentralised governance does not ensure decentralised development. Contributor activity in Lido follows a pattern of structural inequality similar to that observed in traditional open-source projects. To analyse decentralisation more precisely, we propose distinguishing between governance, operational, and developmental decentralisation, as these dimensions do not necessarily align.

Headline contributor counts obscure the concentration of meaningful effort. In Lido, 98.1% of weighted activity originates from the core group. This demonstrates the limits of participation-based metrics and supports the use of weighted activity models. Our approach offers a replicable framework for empirically assessing decentralisation across DeFi and other blockchain ecosystems.

Efforts to reduce centralisation must address knowledge asymmetries and structural access. Modular codebases, rotating code review authority, and mentorship pipelines can facilitate wider participation. Without sustained measurement and institutional support, however, decentralisation may remain more of a rhetorical claim than a practical reality.

6. Threats to validity

While our empirical analysis provides robust evidence for development centralisation in DeFi, several limitations should be acknowledged within the context of our single-case study design.

Single protocol analysis. While single-case designs limit statistical generalisability, Lido represents a best-case scenario given its prominence (largest liquid staking protocol by TVL), maturity (54-month history), and DAO governance structure. If centralisation persists in such favourable conditions, it likely exists across less mature or smaller protocols, making this a particularly informative case for the ecosystem.

GitHub activity as development proxy. Private communications and off-platform coordination may influence development patterns not captured in public repositories. We mitigated this by focusing on Lido's primary repositories, where core protocol implementation occurs and analysing multiple activity types (commits, pull requests, issues, comments) to provide coverage of observable development activities. The consistency of patterns across all activity types strengthens our confidence in the findings.

Activity weighting scheme. Our weighting system (pull requests: 5, commits: 3, issues: 2, comments: 1) reflects qualitative judgments about contribution complexity. However, our extensive sensitivity analysis (Section 3) using Monte Carlo simulation (10,000 iterations) and multiple weighting scenarios demonstrated that concentration patterns persist across all reasonable weight configurations, with Gini coefficients consistently ranging from 0.799 to 0.865. This robustness check validates our methodological choices.

Temporal aggregation. Monthly aggregation may obscure short-term dynamics but provides appropriate balance between statistical stability and temporal resolution. Our 54-month observation period captures multiple development cycles, market crises (Terra/Luna, FTX), and governance decisions, enabling robust identification of both trends and event-driven patterns.

Network construction. Comment-based interaction networks assume co-participation indicates meaningful collaboration. While this assumption may not capture all forms of contributor interaction, the convergence of findings across network metrics (betweenness Gini = 0.854) and activity metrics (weighted Gini = 0.82) suggests that our network construction captures meaningful patterns of centralisation. The consistent identification of hub-and-spoke structures and core–periphery divisions across different analytical approaches strengthens the validity of our conclusions.

Despite these limitations inherent to single-case designs, our multi-method approach, extensive temporal coverage, and convergent evidence across different analytical frameworks provide strong support for our conclusions. The case study methodology enables deep, contextual understanding of development centralisation mechanisms that broader but shallower analyses might miss. Future research should extend this framework to multiple protocols to assess the generalisability of these patterns across the DeFi ecosystem.

7. Conclusions

This study provides a detailed empirical assessment of decentralisation within the open-source development processes of DeFi protocols, using Lido as a critical case study. Despite operating under a decentralised governance framework and maintaining publicly accessible code repositories, our analysis demonstrates that Lido's development activity is highly centralised. A small subset of core contributors accounts for the vast majority of meaningful technical actions, with 98.1% of weighted development efforts originating from the core group.

The application of distributional metrics, such as the Gini coefficient (0.82) and Herfindahl–Hirschman Index, consistently indicates high levels of contributor concentration over time. Moreover, this concentration has increased throughout the protocol's evolution, rising from 0.686 in the bootstrap phase to 0.817 in maturity. Interaction network analysis reveals high centralisation in communication flows (betweenness Gini = 0.854), with key contributors occupying central coordination roles. The fragmentation into 16 disconnected components and the stark core–periphery divide, where peripheral contributors are numerically equal but contribute only 1.9% of activity, reflect structural barriers to broader participation.

These findings align with the concept of the "decentralisation illusion", where protocols branded as decentralised rely on centralised operational structures. This poses risks to governance integrity, protocol sustainability, and resilience against contributor attrition or capture. The 3-month lead time between development activity and TVL changes further suggests that concentrated development control precedes value accrual in the protocol.

Future research should expand this single-case study approach across multiple DeFi protocols to determine whether similar patterns persist throughout the ecosystem. Methodologically, large language models could be used to semantically interpret commit messages and code changes, enabling classification of contributions by type (e.g., security patches, feature additions, refactoring) and assessment of contribution complexity beyond simple activity counts. Such approaches would complement quantitative concentration metrics with qualitative insights into the nature and impact of development work.

Additionally, exploring the impact of incentive mechanisms, governance reforms, and community engagement strategies could offer pathways to fostering more genuinely decentralised development environments. For stakeholders and policymakers, these insights highlight the importance of looking beyond governance tokens and DAOs when assessing the true decentralisation of critical financial infrastructure in the DeFi space.

CRediT authorship contribution statement

Giuseppe Destefanis: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jiahua Xu:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Silvia Bartolucci:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] R. Auer, B. Haslhofer, S. Kitzler, P. Saggese, F. Victor, The technology of decentralized finance (DeFi), Digit. Financ. 6 (1) (2024) 55–95.

[2] M. Aquilina, J. Frost, A. Schrimpf, Decentralized finance (DeFi): a functional approach, J. Financ. Regul. 10 (1) (2024) 1–27.

[3] S. Aufiero, G. Ibba, S. Bartolucci, G. Destefanis, R. Neykova, M. Ortu, The network structure of smart contracts in ethereum dapps, Complex Netw. 2023 (2023).

[4] L. Mungo, S. Bartolucci, L. Alessandretti, Cryptocurrency co-investment network: token returns reflect investment patterns, EPJ Data Sci. 13 (1) (2024) 11.

[5] R. Martins, Web3 in Financial Services: How Blockchain, Digital Assets and Crypto are Disrupting Traditional Finance, Kogan Page Publishers, 2024.

[6] K. Gogol, C. Killer, M. Schlosser, T. Bocek, B. Stiller, C. Tessone, Sok: Decentralized finance (DeFi)–fundamentals, taxonomy and risks, 2024, arXiv preprint arXiv:2404.11281.

[7] G. Ibba, S. Aufiero, R. Neykova, S. Bartolucci, M. Ortu, R. Tonelli, G. Destefanis, A curated solidity smart contracts repository of metrics and vulnerability, in: Proceedings of the 20th International Conference on Predictive Models and Data Analytics in Software Engineering, 2024, pp. 32–41.

[8] G. Ibba, S. Khullar, E. Tesfai, R. Neykova, S. Aufiero, M. Ortu, S. Bartolucci, G. Destefanis, A preliminary analysis of software metrics in decentralised applications, in: Proceedings of the Fifth ACM International Workshop on Blockchain-Enabled Networked Sensor Systems, 2023, pp. 27–33.

[9] G. Ibba, S. Aufiero, S. Bartolucci, R. Neykova, M. Ortu, R. Tonelli, G. Destefanis, MindTheDApp: A toolchain for complex network-driven structural analysis of ethereum-based decentralized applications, IEEE Access 12 (2024) 28382–28394.

[10] G. Ibba, R. Neykova, M. Ortu, R. Tonelli, S. Counsell, G. Destefanis, A machine learning approach to vulnerability detection combining software metrics and topic modelling: Evidence from smart contracts, Mach. Learn. Appl. (2025) 100759.

[11] S. Aufiero, G. Ibba, S. Bartolucci, G. Destefanis, R. Neykova, M. Ortu, Dapps ecosystems: Mapping the network structure of smart contract interactions, EPJ Data Sci. 13 (1) (2024) 60.

[12] M. Aquilina, G. Cornelli, J. Frost, L. Gambacorta, Cryptocurrencies and decentralised finance: functions and financial stability implications, BIS Pap. (2025).

[13] F.S. Board, The financial stability risks of decentralised finance, 2023, https://www.fsb.org/2023/02/the-financial-stability-risks-of-decentralised-finance/. (Accessed 29 December 2023) at.

[14] S. Aufiero, S. Bartolucci, F. Caccioli, P. Vivo, Mapping microscopic and systemic risks in TradFi and DeFi: a literature review, 2025, arXiv preprint arXiv:2508.12007.

[15] M. Dotan, A. Yaish, H.-C. Yin, E. Tsytkin, A. Zohar, The vulnerable nature of decentralized governance in defi, in: Proceedings of the 2023 Workshop on Decentralized Finance and Security, 2023, pp. 25–31.

[16] S. Meneguzzo, C. Schifanella, V. Gatteschi, G. Destefanis, Evaluating DAO sustainability and longevity through on-chain governance metrics, 2025, arXiv preprint arXiv:2504.11341.

[17] J.F. Doerr, A. Kosse, A. Khan, U. Lewrick, B. Mojon, B. Nolens, T. Rice, DeFi risks and the decentralisation illusion, BIS Q. Rev. 21 (2021).

[18] M. Vaccargiu, S. Aufiero, S. Bartolucci, R. Neykova, R. Tonelli, G. Destefanis, Sustainability in blockchain development: A bert-based analysis of ethereum developer discussions, in: Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering, 2024, pp. 381–386.

[19] M. Vaccargiu, S. Aufiero, C. Ba, S. Bartolucci, R. Clegg, D. Graziotin, R. Neykova, R. Tonelli, G. Destefanis, Mining a decade of event impacts on contributor dynamics in ethereum: A longitudinal study, in: 2025 IEEE/ACM 22nd International Conference on Mining Software Repositories, MSR, IEEE Computer Society, 2025, pp. 552–563.

[20] G. Destefanis, S. Bartolucci, D. Graziotin, R. Neykova, M. Ortu, Introducing repository stability, in: Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering, 2025, pp. 555–560.

[21] Y. Luo, Y. Feng, J. Xu, P. Tasca, Piercing the Veil of TVL: DeFi Reappraised, in: International Conference on Financial Cryptography and Data Security, FC, 2025, URL https://fc25.ifca.ai/preproceedings/94.pdf.

[22] L. Finance, Lido finance project, 2025, https://lido.fi/. (Accessed 22 April 2025).

[23] L. Finance, Lido finance GitHub repository, 2025, https://github.com/lidofinance. (Accessed 22 April 2025).

[24] R.K. Yin, Case Study Research: Design and Methods, vol. 5, Sage, 2009.

[25] C. Gini, On the measure of concentration with special reference to income and statistics, colorado college publication, Gen. Ser. 208 (1) (1936).

[26] S.A. Rhoades, The Herfindahl-Hirschman index, Fed. Res. Bull. 79 (1993) 188.

[27] K. Dunbar, D.N. Treku, J. Owusu-Amoako, The decentralization enigma in DeFi: Impact of US federal funds rate changes, Br. Account. Rev. (2025) 101613.

[28] T. Chełkowski, P. Gloor, D. Jemielniak, Inequalities in open source software development: Analysis of contributor's commits in apache software foundation projects, PLoS One 11 (4) (2016) e0152976.

[29] A. Masuda, T. Matsuodani, Knowledge of time-bin data selection using gini index based type classification in GitHub, Procedia Comput. Sci. 207 (2022) 1783–1791.

[30] A. Agrawal, A. Rahman, R. Krishna, A. Sobran, T. Menzies, We don't need another hero? the impact of "heroes" on software development, in: Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice, 2018, pp. 245–253.

[31] A. Mockus, R.T. Fielding, J.D. Herbsleb, Two case studies of open source software development: Apache and Mozilla, ACM Trans. Softw. Eng. Methodol. (TOSEM) 11 (3) (2002) 309–346.

[32] M. Vaccargiu, R. Neykova, N. Novielli, M. Ortu, G. Destefanis, More than code: Technical and emotional dynamics in solidity's development, in: 2025 IEEE/ACM 18th International Conference on Cooperative and Human Aspects of Software Engineering, CHASE, IEEE, 2025, pp. 260–271.

[33] L. Betti, L. Gallo, J. Wachs, F. Battiston, The dynamics of leadership and success in software development teams, Nat. Commun. 16 (1) (2025) 1–11.

[34] M.P. Rombach, M.A. Porter, J.H. Fowler, P.J. Mucha, Core-periphery structure in networks, SIAM J. Appl. Math. 74 (1) (2014) 167–190.

[35] C. Amrit, J. Van Hillegersberg, Exploring the impact of soclo-technlcal core-periphery structures in open source software development, J. Inf. Technol. 25 (2) (2010) 216–229.

[36] S. Bartolucci, F. Caccioli, F. Caravelli, P. Vivo, Ranking influential nodes in networks from aggregate local information, Phys. Rev. Res. 5 (3) (2023) 033123.

[37] M. Newman, Networks, Oxford University Press, 2018.

[38] R.F. Engle, C.W. Granger, Co-integration and error correction: representation, estimation, and testing, Econ.: J. Econ. Soc. (1987) 251–276.

[39] V.D. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, J. Stat. Mech. Theory Exp. 2008 (10) (2008) P10008.

[40] X. Zhang, T. Martin, M.E. Newman, Identification of core-periphery structure in networks, Phys. Rev. E 91 (3) (2015) 032803.