

From Consensus to Multi-Polarity: A Network Analysis of the Structural Transformation in the European Parliament

Nicola Lo Russo¹, Johan Sauser², Dino Carpentras¹, Arthur Rondeau²,
Didier Wernli², Dirk Helbing^{1,3}, Roland Bouffanais^{2,*}

¹ETH Zurich, Switzerland

²University of Geneva, Switzerland

³Complexity Science Hub Vienna, Austria

Abstract

We analyse 7,191 final roll-call votes cast across five European Parliament (EP) legislatures (EP6–EP10, 2004–2026). Across four converging analyses we find a coherent transformation from a consensus-oriented chamber, able to assemble broad and stable majorities, to a fragmented, multi-polar one. First, aggregate indicators document an *agreement crisis*: opposition to the majority and the dispersion of votes rise sharply, while the chamber’s once-dominant coalition pattern gives way to shifting alignments with no clear successor. Second, community detection reveals a roughly fourfold rise in network modularity, a shift from a diffuse consensus network to rigid, well-demarcated voting blocs; as this hardening sets in, formal party labels lose their grip on divisions that increasingly set the political mainstream against its challengers, a split that cuts across group lines and tends to follow a pro-EU versus sovereigntist axis. Third, decomposing modularity by policy area shows the hardening is systemic yet most pronounced in Environment and Health, the domain most closely associated with the European Green Deal and the well-sampled area where the mainstream centre-left and centre-right align least, divided by GMO authorisations and the most contested climate files. Fourth, a multiplex analysis of per-vote coalition geometries reveals that the EP is moving from a small repertoire of well-separated voting contexts to many weakly-separated, co-existing ones. Together these establish EP10 as structurally distinct from the preceding two decades, a configuration whose persistence, and its possible implications for the Parliament’s exposure to outside influence, its ongoing term will test.

Keywords: European Parliament, legislative networks, community detection, network modularity, polarisation, Leiden algorithm, multiplex networks

1 Introduction

The European Parliament (EP) is one of the world’s most complex supranational legislatures, representing over 450 million citizens across 27 member states. Over the past two decades it has been confronted with a sequence of structural shocks (the 2004 and 2007 enlargements, the global financial crisis, the migration crisis, Brexit, the COVID-19 pandemic, and the war in Ukraine), while simultaneously expanding its own legislative authority through the Maastricht (1992), Amsterdam (1997), and Lisbon (2007) Treaties [1, 2]. These shocks unfolded amid a broader politicisation of European integration: the permissive consensus that long allowed elites to advance integration with little public pushback has given way to a constraining dissensus [3], while a new transnational cleavage, driven by challenger parties

and pitting pro-integration forces against nationalist and sovereigntist ones, has reshaped party competition across Europe [4]. Political scientists have documented the resulting changes in seat shares, party group membership, and alliance patterns, using qualitative and descriptive methods [5–8]. Yet quantitative accounts of whether these shifts constitute a temporary or a genuine *topological* transformation of the Parliament’s internal structure remain few and largely confined to single terms or short windows. Research on the U.S. Congress established that legislative polarisation can be measured via network modularity: treating pairwise co-voting as edges and applying community-detection algorithms recovers politically meaningful clusters whose cohesion over time tracks partisan sorting [9, 10]. A large political-science literature has likewise mapped the empirical structure of EP voting: transnational party groups, rather than national delegations, are the primary

*Corresponding author: Roland.Bouffanais@unige.ch

organising units [8, 11]; group cohesion is high and has risen across successive enlargements [7, 12, 13]; competition is structured chiefly by a dominant left–right dimension alongside a secondary pro-/anti-EU axis [8, 14–16]. Network-analytic studies have applied community detection to the EP, but primarily to characterise a single legislature or to compare a small number of terms [17, 18]. Spatial roll-call models have themselves been applied longitudinally to the EP’s earlier terms, tracking the consolidation of the left–right dimension and the rise of group cohesion [5, 19]; we complement that ideal-point tradition with a network-topological lens, which recovers multidimensional coalition geometry without fixing the number of dimensions a priori. To our knowledge, no study has applied such a network-modularity framework longitudinally across five complete or ongoing EP terms (EP6–EP10), spanning the full period from post-enlargement consensus through the current multi-polar configuration.

We fill this gap with four contributions that apply established network and community-detection methods to this longitudinal question. First, we assemble a harmonised dataset of 7,191 main-vote roll-call records across EP6–EP10 (2004–2026) and document an *agreement crisis* through aggregate statistics. Second, we construct a weighted MEP-agreement network per legislature and apply the Leiden community-detection algorithm to track modularity and community composition across time. Third, we disaggregate the modularity analysis by policy area and show that structural hardening is systemic, but is amplified disproportionately by the Environment and Health domain, where GMO authorisations and the most contested Green Deal files divide the mainstream centre-left and centre-right as the EU’s environmental agenda intensified from EP8 onward. Fourth, we apply the multiplex signed-network framework of Arinik et al. [20] to identify distinct voting contexts per legislature and quantify the growing coexistence of multiple coalition logics. Together, these analyses provide a high-resolution empirical portrait of the EP’s transition from consensus to multi-polarity.

Figure 1 previews this transformation schematically: the space of political positions, drawn as a sphere, is cut by a growing number of cleavage planes, from a single dominant axis of conflict in EP6 to fractures in every direction in EP10.

2 Data

2.1 Sources and Coverage

The empirical foundation of this study is a longitudinal dataset spanning five EP terms (four complete, one ongoing) from 2004 to January 2026. We begin the analysis with EP6 (2004–2009), the first full term

following the 2004 enlargement that brought the Parliament to its present scale and is a natural baseline for the modern, post-enlargement EP. Records from EP6 through the first part of EP9 (2004–2022) were sourced from VoteWatch Europe [21], an independent parliamentary monitoring organisation whose historical roll-call archive was retrieved before the organisation ceased operations in 2022. Data from 2019 onwards were obtained from the open HowTheyVote.eu REST API [22], which records per-vote outcomes and individual member positions. Both sources permit research use; redistribution of the underlying records remains subject to the original providers’ terms, while the harmonised dataset and all analysis code are released under the MIT licence (see *Availability of Data and Code*). A three-year overlap (2019–2022) was used to validate source consistency: votes present in both sources during this period were cross-checked and found consistent in outcome and member-level positions, enabling the construction of a single longitudinal stream.

2.2 Vote Selection

The EP records four types of roll-call votes (RCVs). We restrict the analysis to the *main (final) votes* on legislative reports and resolutions, which we define as the definitive roll-call decisions that close the Parliament’s consideration of a file. This restriction is justified on two grounds. Substantively, final votes on legislative files imply direct legal consequences under the Ordinary Legislative Procedure [23] and are less prone to tactical signalling than amendment votes, forcing parties to reveal their legislative preferences. Methodologically, the volume of amendment RCVs grew dramatically from EP6 to EP9, a trend generally attributed to the strategic use of roll-call requests by political groups [24, 25]; including amendment votes would therefore introduce selection bias into a cross-term comparison. Final votes, in contrast, have been recorded by default since July 2009 for legislative proposals, an obligation extended by the 2014 amendment to Rule 166 of the EP’s Rules of Procedure to any single or final vote on a report [6, 26]. The EP6 records predate this rule and their final RCVs were requested by political groups; however, Hix et al. [25] show that this selection has a negligible effect on group cohesion, which supports including the EP6 final votes in the longitudinal comparison.

The scale of this selection is substantial. Across the five terms the EP recorded 45,469 roll-call votes, of which only 7,191 (16%) are main votes (Table 1); the remaining votes are dominated by amendment, split, and separate-paragraph votes, whose volume grew sharply over time, most strikingly in EP9, where final votes account for just 9.4% of the 19,272 recorded roll calls (Fig. 2). Retaining this strate-




Figure 1: Schematic view of the structural transformation. The sphere represents the space of political positions in the Parliament; each plane separates opposing opinions on one contentious policy domain. The number of planes equals the effective number of voting contexts measured in Table 6 (3.2 in EP6, 6.3 in EP8, 8.5 in EP10). In EP6 the planes are nearly parallel: conflict is organised along a single dominant axis, the economic left–right divide. By EP10 the planes cut the sphere in every direction, as domains such as European integration, environmental regulation, migration and the rule of law each open a fracture of their own (Section 4.5). Plane colours carry no meaning; they only distinguish the planes.




Figure 2: Final (main) votes are a minority of all recorded roll-call votes (RCVs) in every term. The total RCV count rises sharply over time, peaking in EP9, where only 9.4% of the 19,272 recorded votes are final, driven by strategically inflated amendment voting. Restricting the analysis to final votes (16% of all RCVs across the period) removes this selection bias. EP10’s lower totals reflect the ongoing term (the first 19 months as of this writing).

gically inflated amendment record in a cross-term comparison would introduce selection bias [24, 27], which the restriction to main votes avoids.

The resulting dataset comprises 7,191 main votes (Table 1). EP10 covers only the first 19 months of the current term as of this writing.

2.3 Voting Positions and Pre-Processing (Filtering)

Each roll-call records one of four positions per MEP: FOR, AGAINST, ABSTENTION, or DID_NOT_VOTE (DNV). Deliberate abstention is recorded separately as the ABSTENTION position and retained. DNV, by contrast, conflates procedural artefacts (mid-term

replacements, seats not yet occupied) with incidental absences such as illness; although non-attendance can occasionally be political, it cannot be read as a directional for/against choice and is dominated by these non-political factors. We therefore exclude DNV from all agreement and entropy calculations throughout.

To ensure statistical reliability and longitudinal comparability, we apply two filters uniformly across all analyses: (i) votes in which a single position (FOR, AGAINST or ABSTENTION) is taken by more than 97.5% of participating MEPs (near-consensus votes) are excluded, as they carry negligible information about ideological cleavages; and (ii) MEPs with participation rates below 50% are excluded from network construction to avoid distortions introduced by short-term replacements or chronic absenteeism.

In practice these filters are mild. The near-consensus criterion removes only 0.6–6.1% of main votes per term (5.8%, 3.9%, 0.6%, 6.1%, and 2.8% for EP6–EP10; the EP9 rebound mirrors the temporary grand-coalition consolidation documented in Section 4.1), because final votes are already substantively contested; the cross-term conclusions are therefore insensitive to the exact threshold, and the resulting non-consensus counts appear in Table 3. The participation filter retains 667–714 active members from the 745–959 recorded per term. Both filters apply identical thresholds in every legislature, so they cannot by construction account for the cross-term trends reported below, which reflect changes in the underlying voting structure rather than the choice of cut-off; the agreement network is

Table 1: Dataset overview. EP10 is ongoing; asterisks denote partial data. *Active MEPs* are those meeting the 50% participation threshold used for network construction (the node counts of Table 3); the gap from recorded MEPs reflects short-term replacements removed by filtering (see Sec. 2.3); the MEPs column is not summed, since members serving in several terms would be counted more than once.

EP	Period	Recorded RCVs	Main votes	MEPs	Active MEPs
EP6	2004–2009	6,199	1,253	940	696
EP7	2009–2014	6,961	1,664	853	704
EP8	2014–2019	10,252	2,072	858	714
EP9	2019–2024	19,272	1,805	870	667
EP10*	2024–	2,785	397	734	710
Total	2004–2026	45,469	7,191	–	–

in any case invariant to the related co-vote overlap threshold (Section 3.1). Among cast positions, FOR predominates throughout, but the AGAINST share roughly doubles, from 9.5% of cast votes in EP6 to 20.9% in EP10. This is an early, albeit purely descriptive signal of the rising contestation quantified in Section 4.1.

3 Methods

Figure 3 summarises the analysis pipeline. Each legislature’s roll-call record yields a set of aggregate contestation indicators directly, and is then processed along two network paths that share the pairwise scoring of Table 2 but differ in their inputs: agreement scores averaged across all shared votes form a weighted MEP-agreement network that feeds community detection (global and per-topic), whereas the same scoring when applied to each vote separately yields per-vote signed networks that feed the multiplex voting-context analysis.

3.1 Pairwise Agreement Scoring

We define a pairwise agreement score between MEPs i and j as follows. Let $\mathcal{V}_{ij}$ denote the set of votes in which both cast a recorded position (FOR, AGAINST, or ABSTENTION). For each vote $v \in \mathcal{V}_{ij}$, a raw score $w(v_i, v_j)$ is assigned according to the matrix in Table 2. The raw similarity is then averaged and normalised:

$$S(i, j) = \frac{1}{|\mathcal{V}_{ij}|} \sum_{v \in \mathcal{V}_{ij}} w(v_i, v_j), \quad (1)$$

$$A(i, j) = \left(\frac{S(i, j) + 1}{2} \right)^2, \quad (2)$$

yielding $A(i, j) \in [0, 1]$, where 0 indicates systematic opposition and 1 indicates perfect alignment. The value 0.25 corresponds to uncorrelated voting. MEP pairs sharing 20 or fewer co-votes are assigned the neutral score $S = 0$ (equivalently $A = 0.25$) to prevent unstable estimates from thin overlaps.

This rule is a safeguard rather than a binding filter: under the 50% participation requirement, no MEP pair shares fewer than 53 votes in any legislature (median overlap 311–1,541), so the global agreement matrices, and every result derived from them, are invariant to the threshold across the range $[10, 50]$ ($Q_{\max}$ identical to four decimals).¹

Table 2: Agreement scoring matrix. DNV excluded from all computations.

	FOR	AGAINST	ABSTAIN
FOR	+1	−1	−0.25
AGAINST	−1	+1	−0.25
ABSTAIN	−0.25	−0.25	+0.25

The treatment of abstentions deserves explicit justification. Abstention is a deliberate choice to withhold a substantive position, distinct from absence; it should therefore influence the similarity score but with reduced weight compared to a direct agreement or disagreement. The ± 0.25 factor is chosen such that four shared abstentions carry the same weight as one shared FOR/AGAINST vote, reflecting partial but genuine alignment [18, 28]. Its sign follows the direction of the alignment: two MEPs who both abstain (a *shared* abstention) score +0.25 (weak agreement), whereas an abstention paired with a FOR or an AGAINST vote (a *mixed* abstention) scores −0.25 (weak disagreement). The precise value is not critical: Section 5 shows the conclusions are unchanged under alternative abstention weights, including a weight of zero, which treats abstentions as carrying no information at all. The quadratic normalisation in Eq. (2) amplifies separation between high-agreement and moderate-agreement pairs, while preserving rank order, thereby improving the resolution of the community-detection step. This nonlinear step is our deliberate choice rather than a standard

¹The threshold can bind only in the sparser per-topic networks, where a pair may share few votes within a single domain; those analyses carry their own explicit guards (topic-legislature pairs with fewer than 20 non-consensus votes or 10 active MEPs are excluded) and their estimates are flagged as noisier.




Figure 3: Analysis pipeline. Per-legislature roll-call records yield aggregate contestation indicators directly (top). The records then feed two network analyses with distinct inputs: pairwise agreement scores averaged over all shared votes form a single weighted MEP-agreement network, the input to modularity-based community detection (global and per-topic, upper right); the same Table 2 scoring applied to each vote separately yields a stack of per-vote signed networks, the input to the multiplex voting-context decomposition (lower right).

encoding from the literature; robustness with respect to it and to the other encoding choices is examined in Section 5, where the linear alternative $A = (S+1)/2$ reproduces the same cross-term trend.

3.2 Network Construction

For each legislature, we construct a weighted, undirected network $G = (V, E, w)$ in which each active MEP is a node, and an edge with weight $A(i, j)$ connects every pair with $A(i, j) \geq \tau$. We set $\tau = 0.1$ throughout. Below this threshold, edge weights are close to zero and reflect noise rather than structural alignment. Retaining them would inflate the null-model baseline in the modularity computation, suppressing Q without contributing to community interpretability. We verify that the qualitative conclusions are stable across the range $\tau \in [0.05, 0.20]$.

The resulting networks contain 667–714 nodes and 213,000–241,000 edges per legislature (Table 3), with density exceeding 0.97 in all cases; community structure therefore emerges from variation in edge weights rather than from topological disconnection.

Table 3: Network properties per legislature ($\tau = 0.1$).

EP	Nodes	Edges	Non-cons. votes
6	696	236,289	1,180
7	704	237,779	1,599
8	714	235,000	2,059
9	667	213,221	1,694
10	710	241,837	386

3.3 Community Detection and Modularity

We apply the Leiden algorithm [29] to maximise the resolution-parameter (γ) generalisation of network modularity:

$$Q(\gamma) = \frac{1}{2m} \sum_{ij} \left[A_{ij} - \gamma \frac{k_i k_j}{2m} \right] \delta(c_i, c_j), \quad (3)$$

where A_{ij} is the adjacency matrix, $m = \frac{1}{2} \sum_{ij} A_{ij}$ is the total edge weight, $k_i = \sum_j A_{ij}$ is the weighted degree of node i , and $\delta(c_i, c_j)$ equals 1 if nodes i and j share a community. At $\gamma = 1$, this reduces to the standard Newman–Girvan modularity [30]. The reported $Q_{\max}$ values are the standard ($\gamma = 1$) modularity of the optimised partition, placing them on the same scale as the normative-partition scores in Table 4. Because the Leiden algorithm is stochastic, we run it with 20 independent restarts (different seed) per network and report the mean modularity of the optimised partitions, which we denote by $Q_{\max}$; the per-restart spread is negligible (< 0.001), so the mean and the per-run maximum agree up to three decimals.

Two resolution scales are analysed. A *baseline* at $\gamma = 1.00$ yields the natural community structure of the network. A *refined* resolution, selected per legislature by an exhaustive grid search (step 0.01) rather than by fitting, to yield 8–10 communities matching the number of EP party groups, provides a meso-scale view that enables direct comparison with formal institutional structure. This refined resolution is an interpretive lens on community composition rather than the primary evidence for temporal change, which rests on the $Q_{\max}$ trajectory at the natural resolution $\gamma = 1$. The selected values are $\gamma = 1.03$ (EP6), 1.04 (EP7), 1.07 (EP8), 1.06

(EP9), and 1.13 (EP10). We additionally compare $Q_{\max}$ against the modularity achieved by four normative partitions: party affiliation, left–right ideology, mainstream–challenger alignment, and country of origin. The left–right partition places GUE/NGL, Greens/EFA, and S&D on the left and all other families, including the liberals, on the right. For the mainstream–challenger partition [31], the mainstream side holds the EPP, S&D, the liberals (ALDE, later Renew), and the conservatives (ECR); the challenger side holds all remaining families: the radical left (GUE/NGL), the Greens/EFA, the sovereigntist right (UEN, IND/DEM, EFDD, ENF, ID, P/E, ESN), and the non-attached. This split is close to, but not identical with, the traditional grand coalition that has dominated EU institutions since their inception (the ECR sits on the mainstream side). It is important to note that is not a pure pro-/anti-EU divide (the pro-integration Greens count among the challengers). The country partition serves as a structural null, since the EP is not organised along national lines. When reporting community composition, each detected community is labelled by its dominant (plurality) party group; we report the share of members belonging to that group (its *purity*) together with the three largest constituent groups, and composite labels (e.g., “grand coalition”) denote communities in which several mainstream groups jointly constitute the membership.

To validate the Leiden results and confirm they are not an artefact of modularity optimisation, we replicate the community-detection step using two spectral methods derived from the stochastic block model (SBM). The first is adjacency spectral embedding with Gaussian Mixture Model clustering (spectral SBM). The second is Laplacian spectral embedding with unit-normalised rows, a degree-corrected variant known as the degree-corrected SBM (DCSBM). Both use BIC-based model selection [32, 33].

3.4 Per-Topic Modularity

The global modularity analysis aggregates all votes regardless of policy area. To identify which domains drive this structural hardening (the strengthening of voting-bloc separation reflected by rising modularity), we repeat the full pipeline – agreement matrix computation, network construction, Leiden clustering – independently for each of the 20 policy areas corresponding to EP standing committees (e.g., Foreign Affairs, Environment and Health, Economic & Monetary Affairs). Topic-legislature pairs with fewer than 20 non-consensus votes or fewer than 10 active MEPs are excluded to avoid spurious modularity estimates on sparse networks.

For each topic pair within the same legislature, we also compute the Pearson correlation r between the upper triangles of the two per-topic MEP agree-

ment matrices, restricted to MEPs who participated in both topics. This comparison uses the raw mean-agreement scores $S(i, j)$ of Eq. (1), prior to the quadratic normalisation, because it concerns the co-variation of agreement patterns across domains rather than their use as community-detection weights. A value of $r = 1$ indicates that the same MEP pairs agree (or disagree) across both domains; $r = 0$ indicates that the two domains activate independent cleavages; and $r < 0$ would indicate domains that activate opposing alignments (in practice, however, only a few near-zero negative pairs occur, all in EP10; Section 4.4). Averaging r across all pairs within a legislature yields a summary measure of cross-topic coherence.

3.5 Multiplex Voting Contexts

The analyses above treat each legislature as a single static network. The multiplex approach of Arinik et al. [20], developed to recover the distinct coalition structures latent in signed voting networks, asks instead: do different votes produce different coalition patterns, and can we group votes by the structure they generate?

For each roll-call vote v , we construct a per-vote signed network G_v whose edge weights apply the same pairwise scoring as the aggregate network (Table 2): +1 for agreement, −1 for disagreement, +0.25 for a shared abstention (both MEPs abstain), and −0.25 for a mixed abstention (one abstains while the other votes FOR or AGAINST). Whereas the aggregate agreement score (Eq. (2)) averages these values over all shared votes, the multiplex construction retains the signed per-vote weights, so that each vote yields its own coalition geometry. Each G_v is partitioned by applying the Leiden algorithm to the positive-edge subgraph, yielding a partition π_v . Disagreement edges therefore determine which MEPs are *not* positively linked but do not themselves drive the partition; this step is thus a per-vote positive-agreement clustering rather than a full signed (frustration-based) community detection. We adopt this simplification because full signed optimisation is prone to numerical instability and produces degenerate, trivial partitions on the sparse single-vote graphs considered here. We compare all partition pairs within a legislature using Normalised Mutual Information (NMI):

$$\text{NMI}(\pi_i, \pi_j) = \frac{I(\pi_i; \pi_j)}{\sqrt{H(\pi_i) \cdot H(\pi_j)}}, \quad (4)$$

restricted to the MEPs who participated in both votes; pairs sharing fewer than ten common members are left undefined and imputed at the mean similarity. This imputation is non-binding in practice: the participation and non-consensus filters leave every pair of votes with a large shared electorate (median

455–570 common MEPs across the five terms), so only three of the roughly 5.6 million within-term vote pairs in EP6–EP10 (all in EP7, minimum overlap nine voters) fall below the ten-member threshold, and the mean substitution has no material effect on the similarity matrix. To cluster the resulting $n \times n$ similarity matrix $\mathbf{S}$ into voting contexts, we first remove the grand-coalition baseline. We eigen-decompose $\mathbf{S}$ and discard the dominant eigenvector u_1 , which captures the consensus signal shared by all votes. Each vote is then embedded in the residual eigenvectors $u_2, u_3, \dots$, retained while their eigenvalue remains above 10% of λ_2 , with each axis scaled by $\sqrt{\lambda}$. Euclidean distances in this residual embedding are clustered with Ward’s linkage, and the number of contexts k is selected by maximising the silhouette score over $k \in \{2, \dots, 10\}$. Each resulting cluster defines a *voting context*: a group of roll calls producing similar coalition geometries. We validate context stability against a random-label null model (200 permutations), bootstrap resampling (50 runs, 80% of votes), and, for EP10, a 5%-label-noise perturbation test. The 10%-of- λ_2 eigenvalue-retention threshold is likewise not load-bearing. Sweeping it across 5%–20% varies the number of retained eigenvectors several-fold. The optimal context count nonetheless remains stable to within two (identical for EP6 and EP9), and the qualitative growth from about five contexts in EP6–EP7 to eight–ten in EP8–EP10 is preserved. The recovered context partitions remain moderately to strongly correlated with the 10% solution (adjusted Rand index 0.58–0.84, weakest for EP8 and EP10).

4 Results

4.1 Aggregate Indicators: The Agreement Crisis

We characterise the aggregate voting dynamics of each legislature through three complementary measures: the opposition-to-majority share, vote entropy, and coalition-pattern diversity.

Opposition-to-majority share. For a vote with n_f, n_a, n_s participants voting FOR, AGAINST, and ABSTAIN respectively, the opposition share is:

$$\text{Opp} = 1 - \frac{\max(n_f, n_a, n_s)}{n_f + n_a + n_s}. \quad (5)$$

The distribution of this quantity shifts markedly rightward across legislatures. The median rises from 0.095 in EP6 to 0.193 in EP8; after a dip to 0.159 in EP9, it reaches 0.252 in EP10. EP8 marks the first clear break from the low-contestation regime of EP6–7 (Fig. 4).

Shannon entropy. Vote entropy $H = -\sum_k p_k \log_2 p_k$ (summed over FOR, AGAINST, ABSTAIN proportions) is bounded in $[0, \log_2 3] \approx [0, 1.58]$ bits. Expressed as a fraction of this maximum, mean entropy rises from 36.8% (EP6) to 40.6% (EP7), 52.7% (EP8), 48.4% (EP9), and 57.1% (EP10). The EP9 dip reflects a temporary consolidation of the grand coalition under the pressure of Brexit and COVID-19; the EP10 level represents the historical maximum for the full period. Chronological tracking within each legislature confirms that EP10 is an outlier from its opening sessions, ruling out a start-of-term effect (Fig. 4).

Coalition-pattern diversity. We define a coalition pattern as the string of per-party majority positions (FOR, AGAINST, ABSTENTION, or tie) for each roll-call vote. The Herfindahl–Hirschman Index ($\text{HHI} = \sum_i p_i^2$, where p_i is the share of votes matching pattern i , measures concentration. Lower HHI indicates more distributed coalition variety.

EP9 is the most concentrated ($\text{HHI} = 0.076$), driven by a dominant all-groups-FOR pattern accounting for 20.4% of all votes. EP10 is the least concentrated ($\text{HHI} = 0.022$): no pattern exceeds 8%, and the top-five patterns collectively account for only 23% of all votes, compared to 49% in EP9. This collapse in coalition predictability is the most direct evidence that the EP no longer has a stable standing legislative majority: the composition of the winning coalition has become substantially less predictable, shifting from one vote to the next.

Figure 6 shows the full repertoire of patterns for each legislature: every tile is one distinct coalition pattern, and its area reflects the share of votes it accounts for. Because the number of distinct patterns grows mechanically with the length of the vote record, we control for record length by down-sampling every legislature to EP10’s 397 main votes (200 draws). The effective number of patterns, measured as the inverse Simpson index of the pattern-share distribution, is then 19 (EP6), 21 (EP7), 33 (EP8) and 13 (EP9), against 50 in EP10: at an identical vote count EP10 sustains close to four times the coalition diversity of EP9 and one and a half times that of EP8.

The leading patterns also change in kind: the most frequent configuration in which the EPP and S&D take opposite positions ranks 17th in EP6 and 8th in EP9, but second in EP10, where it alone accounts for 4.0% of all votes. Configurations that divide the grand coalition are no longer a structural rarity.

Three structural trends are notable in the pattern data. First, the strict all-FOR pattern falls from 18.7% (in EP7) and 20.4% (in EP9) to 7.1% (in EP10). Second, the center-versus-right-wing-




Figure 4: Three complementary measures of the agreement crisis across EP6–EP10. *Left*: Per-vote opposition-to-majority share (violin plots; bars mark the quartiles); the median rises from 0.095 (EP6) to 0.252 (EP10), dipping to 0.159 in EP9. *Centre*: Coalition-pattern concentration curve (cumulative share of votes against the ranked share of coalition patterns); EP10 is the least concentrated legislature on record (HHI = 0.022). *Right*: Per-vote vote entropy as a percentage of its maximum (violin plots); the distribution shifts upward across legislatures, highest in EP10. All three panels are computed over all main votes; the structural analyses of later sections use only the non-consensus subset.




Figure 5: Distribution of pairwise MEP agreement scores per legislature. The distribution shifts from concentrated at high agreement (EP6–7, medians near 0.80) to near-symmetric and bimodal (EP10, mean and median near 0.51), with the share of MEP pairs below the midpoint ($A < 0.5$) rising from about 21% to nearly half.

opposition family of patterns (GUE/NGL through the EPP FOR, with at least one right-wing group AGAINST) peaks at 36.6% of votes in EP8, when the mainstream groups still largely refused cooperation with the far right and routinely voted together against it (the practice known as the *cordon sanitaire*), before declining to 28.0% in EP10, where the boundary between “centre” and “right” itself becomes unstable. Third, cross-bloc voting (parts of the left opposing while parts of the right support) surges from 18.6% (EP9) to 35.3% of EP10 votes.

Agreement score distributions. Applying Eq. (2) to all MEP pairs within each legislature yields the distribution shown in Fig. 5. In EP6–7 the distributions are concentrated at high agreement, with medians near 0.80 (0.80 and 0.79) and a long left tail, reflecting a chamber in which most pairs agreed most of the time. By EP8 the median falls to 0.60 (mean 0.55) and the distribution broadens.

EP10 presents a near-symmetric, bimodal distribution with mean and median near the theoretical midpoint (0.51 and 0.52), with the share of MEP pairs falling below the midpoint ($A < 0.5$) rising from about 21% in EP6 to nearly half in EP10. This bimodality is statistically robust. A Hartigan dip test rejects unimodality for every legislature (partly a consequence of the very large number of MEP pairs), but EP10’s dip is the largest of the five ($D = 0.025$, $p < 0.001$ by bootstrap). What singles out EP10 is the *balance* of its two modes. A two-component Gaussian mixture is preferred over a unimodal fit by a decisive BIC margin. It centres the two modes at agreement scores of 0.28 and 0.72, astride the midpoint, with near-equal weights (0.48 and 0.52). The upper mode marks pairs that vote closely together; the lower mode sits just above the uncorrelated-voting value of 0.25, therefore it captures not a second disciplined bloc in steady opposition but a large mass of MEP pairs voting




Figure 6: The repertoire of coalition patterns across EP6–EP10. Each rectangle represents all main votes of a legislature; every tile is one distinct coalition pattern, defined as in this section by the string of per-group majority positions, with tile area proportional to the share of votes that pattern accounts for. Tiles are coloured by the relationship between the two poles of the grand coalition: blue where the EPP and S&D take the same position, with darker shades marking the patterns in which more of the other groups join them, and red where they take different positions. In EP6 two patterns alone account for 27% of all votes and the pane is dominated by a few large tiles; EP7 is similar. EP8 begins to break up. EP9 re-concentrates, its largest pattern reaching 20.4%, the highest of the period, consistent with the temporary consolidation of the grand coalition. EP10 then breaks the pattern: no configuration reaches 8%, the second largest tile is one in which the EPP and S&D oppose each other, and the pane dissolves into a fine grain of one-off configurations.

nearly independently of one another. The curve thus records fragmentation rather than a clean two-camp split, consistent with the tri-polar community structure of Section 4.2, whose Eurosceptic pole is by far the loosest internally, and with the many weakly-separated coalition geometries of the multiplex analysis (Section 4.5).

4.2 Network Modularity and Community Structure

Global modularity evolution. Maximum modularity $Q_{\max}$, obtained by Leiden optimisation, increases nearly fourfold from EP6 ($Q = 0.026$) to EP10 ($Q = 0.100$), marking a structural transition from a diffuse consensus network to one characterised by rigid, well-demarcated voting blocs (Table 4; Fig. 7). Because raw modularity is not directly comparable across networks of differing size or density, we confirm this increase against both a vote-count-matched down-sampling control and a size- and density-preserving null normalisation (Section 5): the normalised modularity $Q_{\max}/Q_{\text{null}}$ rises from 2.7 to 6.4, so the trend reflects genuine structural change rather than the smaller vote record of EP10.

The country partition ($Q_{\text{cty}}/Q_{\max} \approx 0.05\text{--}0.08$) is negligible throughout, confirming that the EP does not exhibit a voting function that follows national lines. Two structural transitions are visible in the remaining ratios. In EP6–7, party affiliation is the dominant explanatory dimension, capturing 74–80% of $Q_{\max}$. Between EP7 and EP8, a gap opens: $Q_{\max}$ increases by 66% but the party

ratio drops to 68%, implying that the network’s actual divisions have begun to outpace party labels. By EP10, the mainstream–challenger partition ($Q_{\text{MC}}/Q_{\max} = 0.615$) surpasses the party partition (0.571) for the first time. The two are close, so formal party groups still capture much of the structure; but the crossover signals that the single split between the political mainstream and its challengers now organises voting at least as well as the full set of party-group labels. This divide tends to follow the pro-EU versus sovereigntist axis without reducing to it, since the radical left and the Greens sit on the challenger side and the ECR on the mainstream side.

Spectral SBM and DCSBM community detection recovers structure broadly consistent with the Leiden partitions, and method-agnostic quality metrics (coverage, conductance) favour the Leiden partitions in every legislature, confirming that the community structure is not an artefact of modularity optimisation.

Baseline community structure. At $\gamma = 1.00$, the Leiden algorithm identifies exactly three communities in every legislature. The macro-scale structure is stable in form but has shifted fundamentally in content (Table 5).

The structural contrast between EP8 and EP10 is the most striking. In EP8, a single super-community absorbed 63% of MEPs (EPP + S&D + ALDE), quantifying the institutional centre’s cohesive response to the migration crisis and Brexit. By EP10, the parliament has transitioned to a tri-polar configuration: the Eurosceptic right, which formed a




Figure 7: Evolution of network modularity across EP6–EP10. $Q_{\max}$ (Leiden, filled circles) increases fourfold from 0.026 to 0.100. Normative partition scores are shown for comparison. The country partition (open squares) is negligible throughout.

Table 4: Modularity scores and normative partition ratios. $Q_{\max}$ is the mean Leiden modularity over 20 restarts; the per-restart spread is negligible (< 0.001 in every term, the optimiser being effectively deterministic), so the reported 95% confidence intervals are instead obtained by vote-level bootstrap (300 resamples of each legislature’s votes), capturing sampling rather than optimisation uncertainty. Ratios show what fraction of $Q_{\max}$ each normative partition achieves.

EP	$Q_{\max}$	95% CI	$Q_{\text{party}}/Q_{\max}$	$Q_{\text{LR}}/Q_{\max}$	$Q_{\text{MC}}/Q_{\max}$	$Q_{\text{cty}}/Q_{\max}$
6	0.026	[0.023, 0.029]	0.741	0.714	0.494	0.055
7	0.033	[0.030, 0.036]	0.802	0.652	0.467	0.082
8	0.054	[0.052, 0.057]	0.680	0.546	0.597	0.063
9	0.053	[0.049, 0.057]	0.621	0.621	0.474	0.064
10	0.100	[0.088, 0.112]	0.571	0.554	0.615	0.052

Table 5: Baseline three-community structure ($\gamma = 1.00$). Dominant party and approximate size shown.

EP	C_0	C_1	C_2
EP6	S&D+ALDE (41%)	EPP (36%)	Mixed (23%)
EP7	EPP (37%)	S&D+ALDE (37%)	Mixed (27%)
EP8	Grand coalition (63%)	Eurosceptic (21%)	Greens/GUE (16%)
EP9	S&D+Greens+Renew (47%)	EPP+Renew (30%)	Eurosceptic (23%)
EP10	S&D+Greens (38%)	EPP+Renew (32%)	Eurosceptic (30%)

marginal fringe in EP6, has grown into the third bloc and approaches parity with the two traditional mainstream communities. The Eurosceptic community accounts for 30% of active MEPs in EP10 at baseline resolution.

Refined community structure. At resolution parameters $\gamma \in [1.03, 1.13]$ selected to yield 8–10 communities (Fig. 8), a consistent historical pattern (EPP and S&D as the two largest communities) breaks for the first time in EP10. In every legislature from EP6 through EP9, EPP and S&D form the two

largest refined communities. In EP10 ($\gamma = 1.13$), the Eurosceptic bloc becomes the largest community (214 MEPs, dominated by Patriots for Europe [PfE] 84, European Conservatives and Reformists [ECR] 78, Europe of Sovereign Nations [ESN] 27, and non-attached [NI] 24), exceeding both the EPP core (168 MEPs) and S&D core (156 MEPs). Even excluding the 24 non-attached members (a residual administrative category), the sovereigntist groups proper (PfE, ECR, ESN; 189 MEPs) remain the largest grouping. This “bloc” is, however, an analyst-level aggregate of formally distinct and often rival groups; its internal cohesion is correspondingly low (below), and it is best read as a numerically large oppositional pole that, unlike the EPP or S&D cores, inherently lacks programmatic coherence.

This structural inversion is accompanied by two additional changes. First, EPP fragments across three sub-communities that seem to be separated by genetically-modified-organism (GMO) policy, rule-of-law stances, and trade alignment, the first EPP fragmentation detected in 20 years of data. Second, the liberal group (ALDE, the Alliance of Liberals and Democrats for Europe, in EP6–EP8, reconstituted as Renew Europe in 2019) formed independent communities of 60–90 MEPs through EP8 but, by EP10, no longer resolves as a distinct community at the tested resolutions: its members now cluster within the EPP orbit (40 MEPs, Northern/Western liberals) or the S&D orbit (27 MEPs, social-liberals), leaving only a 7-MEP micro-community. While Renew remains a formally constituted group, its voting behaviour has converged sufficiently with the mainstream blocs that community detection no longer isolates it as a computationally distinct voting force. This reflects convergence: Renew’s own internal cohesion remains high (mean internal edge weight $\bar{w}_{in} = 0.82$, comparable to the EPP and S&D cores), so this convergence indicates that it now votes *with* the grand-coalition centre while retaining its internal discipline.

The Eurosceptic bloc, while numerically the largest, exhibits the lowest internal cohesion ratio of all four main communities ($CR = 1.57$ vs 1.70 for EPP core). The cohesion ratio (CR) is defined as the mean agreement weight among a community’s members, divided by the mean weight of their edges to members of other communities; a value of 1 denotes a community no more internally aligned than with the rest of the chamber. The Eurosceptic bloc’s mean internal edge weight is $\bar{w}_{in} = 0.462$, substantially lower than the EPP (0.855) and S&D (0.891) cores. The bloc’s cohesion comes from systematic opposition to the pro-European majority, with little shared programmatic alignment, consistent with the distinct ideological positions of its constituent party families.

4.3 Per-Topic Modularity: Systemic Hardening and the Environmental Outlier

Figure 9 shows the Q_{max} matrix across 20 policy areas and five legislatures. Two structural patterns emerge.

Universal increase. Every policy area for which sufficient data exists shows higher modularity in EP10 than in EP6. Of the six EP10 domains with enough non-consensus votes (≥ 20), all six exceed their EP6 baseline values, and three exceed the parliament-wide Q_{max} of 0.100. This universality establishes that structural hardening is *systemic*, it affects the entire legislative agenda, rather than being an artefact driven by a single dominant domain.

Environment as the outlier. Despite the systemic trend, one policy area diverges sharply in both rate and magnitude. Environment and Health undergoes a modularity increase from $Q_{max} = 0.012$ (EP6) to $Q_{max} = 0.211$ (EP10), more than doubling the parliament-wide average and representing the largest topic-level modularity recorded in any well-sampled domain across any legislature (Fig. 10). As we show below, the health files remain cohesive; the rise is carried by the environmental files, above all GMO and climate.

Crucially, this acceleration is not contemporaneous with the general trend. Environment and Health, which is indistinguishable from the low-modularity baseline in EP6 and EP7, first surpasses the global Q_{max} in EP8, coinciding both with the contested stream of GMO-authorisation objections that formed the domain’s earliest EPP–S&D fault line and with the wave of climate legislation that preceded the European Green Deal (the 2016 ratification of the Paris Agreement, the 2017–2018 revision of the Emissions Trading System, and the Clean Energy for All Europeans package). It accelerates further through EP9 and EP10, as the Green Deal itself (launched in December 2019, at the start of EP9) advances into legislation. No other policy area presents a comparable trajectory in either timing or magnitude.

The party-label fit (the ratio Q_{party}/Q_{max} for the Environment network) declines monotonically from 59% (EP6) to 44% (EP10), i.e. as Environmental modularity surged, formal party labels progressively lost their ability to explain the resulting network structure. This is precisely what would be expected if the Environment cleavage cuts *within* the traditional pro-EU coalition (EPP versus S&D and Greens) rather than between the mainstream and the Eurosceptic opposition: party groups that vote together on most legislation fracture, however, on GMO authorisations, deforestation regulation, and




Figure 8: Community composition at refined resolution ($\gamma \in [1.03, 1.13]$) for EP6–EP10. In EP6–EP9 the EPP and S&D form the two largest communities. In EP10 the Eurosceptic bloc (Pfe + ECR + ESN + NI) becomes the largest for the first time. Bar segments are coloured by party group.

corporate sustainability due diligence, thereby producing community structures that cross party lines.

The inter-community agreement data confirm this interpretation. Environment and Health shows by far the weakest EPP–S&D alignment of any well-sampled domain: their mean pairwise agreement across the 54 Environment votes collapses to

+0.22, far below the +0.69 chamber-wide level and the +0.80 (Foreign Affairs), +0.73 (Economic and Monetary Affairs), and +0.90 (Budgetary Control) recorded elsewhere. On the agreement scale, where 0 denotes uncorrelated voting, +0.22 is only weakly positive: the grand coalition’s alignment nearly dissolves on Environment, though on average it stops




Figure 9: Per-topic Q_{max} across the 20 policy areas and five legislatures. Policy areas are ordered from the bottom up by their EP10 Q_{max} ; cells marked *n/a* have too few non-consensus votes (< 20) for a reliable estimate, which in the ongoing EP10 term leaves only six policy areas estimable (these occupy the lower rows – the remaining areas appearing above). Most areas show rising modularity through EP6–EP9, and Environment and Health is the structural outlier, rising from 0.012 to 0.211.

just short of net opposition. The collapse is concentrated on two fronts. GMO authorisations are the only Environment sub-area marked by net disagreement (mean EPP–S&D agreement -0.30 across 19

votes), where the fault line becomes visible in EP8–EP9; the most contested Green Deal files produce the sharpest individual splits, with EPP and S&D landing on opposite sides of the European Climate Law and the Deforestation Regulation at agreement scores near -0.95 . Health-related votes, by contrast, remain cohesive ($+0.66$).² The modularity surge itself is broad-based across the Environment agenda rather than GMO-specific: excluding all GMO votes lowers the Environment $Q_{\max}$ only marginally, from 0.21 to 0.19.

4.4 Cross-Topic Agreement Coherence

For each pair of policy areas within the same legislature, we compute the Pearson correlation r between the upper triangles of their MEP agreement matrices, restricted to MEPs who participated in both domains (Section 3.4). A high r indicates that the same MEP pairs that agree on one topic tend to agree on another, a symptom of a low-dimensional political space in which a single dominant cleavage structures most votes regardless of domain.

Mean cross-topic coherence $\bar{r}$ remains high through the middle legislatures, 0.656 (EP6), 0.730 (EP7), and 0.727 (EP8), with the EP7 peak coinciding with the Eurozone crisis, which may have consolidated cross-domain voting through systemic pressure; it then declines through EP9 (0.636) to 0.493 in EP10 (Fig. 11). The decline from its EP7 peak to EP10 is overall steady and represents a near-30% reduction in cross-domain coherence. Because the $\binom{n}{2}$ entries of each per-topic agreement matrix are not independent, every MEP appears in $n - 1$ pairs, we assess significance with a Mantel permutation test that jointly permutes MEP labels (499 permutations) rather than relying on the ordinary Pearson p -value, which is invalid for non-independent matrix entries. Every cross-topic correlation in EP6–EP9, and 97% of those in EP10, is significant at $p < 0.01$, confirming that the reported coherence values reflect genuine cross-topic structure. Because the Pearson coefficient assumes a linear relationship, we also recompute every coherence value with the rank-based Spearman and Kendall coefficients (Appendix B): the declining trend and EP10’s position as the least coherent legislature survive under both (Spearman $\bar{r}$ falls from 0.71 in EP8 to 0.50 in EP10; Kendall $\bar{\tau}$ from 0.54 to 0.37), and per-pair Pearson and Spearman values agree to $\rho \approx 0.92$, so the matrix-to-matrix relationship is monotonic and the linearity assumption does not drive the result.

²Sub-areas within Environment and Health are identified by keyword rules applied to vote titles (health, medical, pharmaceutical, tobacco, and food-safety patterns for the health files; analogous patterns for GMO and climate), with GMO taking precedence where a title matches several patterns.

In EP10, Environment is a structural outlier: its mean correlation with other domains (≈ 0.38) is among the lowest in the chamber (only the sparse Legal Affairs and Petitions domains are lower), against a cross-topic average of 0.493, with individual values ranging from 0.233 (Legal Affairs) to 0.516 (Employment and Social Affairs). That Environment’s agreement structure is the least predictable from the rest of the agenda, with its weakest links to Legal Affairs ($r = 0.233$), Fisheries ($r = 0.239$), and Petitions ($r = 0.259$), is the topological signature of the European Green Deal introducing a fault line only weakly aligned with the parliament’s pre-existing cleavages. Legal Affairs in particular decouples from the agenda, producing the only negative inter-topic correlations in EP10 (with Culture and Education, $r = -0.04$; Transport; and Budgetary Control); these are near-zero and non-significant under the Mantel test, but no negative pairs arise in any earlier legislature.

4.5 Multiplex Voting Contexts

The multiplex analysis reveals the *mechanism* behind the modularity increase: a progressive accumulation of distinct coalition logics, each operating simultaneously.

Table 6 summarises the key results. The silhouette criterion identifies a small, well-separated set of contexts in EP6 and EP7 (five and six); these are genuine interior optima, unchanged when the search ceiling is raised to $k = 40$. In EP8–EP10, by contrast, it no longer isolates a clear optimum below the ceiling: the silhouette is low (0.21–0.26) and nearly flat across k , and extending the search to $k = 40$ shifts its argmax up to $k = 21$ (EP8), 16 (EP9), and 19 (EP10) without producing a sharp peak (Appendix C). The reported k^* for these three terms are therefore lower bounds, not point estimates, and the precise context count in the contentious later terms is *not* identified by this criterion; these legislatures are best described as fragmenting into *many* weakly-separated coalition geometries rather than a fixed number. The qualitative conclusion is nonetheless robust under measures that require no choice of k : the dominant-context share collapses from approximately 50% (EP6) to 19% (EP10), and the effective number of coalition geometries (the inverse Simpson index of the context-size distribution) rises monotonically from 3.2 to 8.5 across the same span. Consistent with the weak separation, ARI yields fewer contexts than NMI, indicating that many NMI-identified contexts share soft boundaries.

Inspection of the party composition within EP10’s ten contexts (Fig. 12) reveals four qualitatively distinct coalition logics operating simultaneously:

1. **Broad consensus** (140 votes, 36%). A single bloc containing all major groups; only ESN dis-




Figure 10: EP10 per-topic $Q_{\max}$ for the six policy areas with sufficient data (≥ 20 non-consensus votes), compared to the global $Q_{\max}$ (dashed line at 0.100); topics exceeding it are shown in red. Environment and Health dominates at 0.211, more than double the parliament-wide value, with Economic & Monetary Affairs (0.126) and Budgets (0.109) also above the global level. Only 6 of 20 policy areas have enough EP10 votes for reliable estimation.

Cross-topic agreement matrix correlation (EP6--EP10)




Figure 11: Cross-topic agreement matrix correlations for EP6–EP10. Each cell shows the Pearson correlation between the upper triangles of two per-topic MEP agreement matrices. Mean cross-topic coherence peaks at $\bar{r} = 0.730$ (EP7) and declines sharply to $\bar{r} = 0.493$ (EP10), with Environment emerging as a persistent structural outlier.

sents. Confined primarily to budget and uncontroversial foreign-affairs votes. Note that this context is identified *within* the non-consensus vote set, since near-unanimous votes (a single position above 97.5% of the votes cast) are al-

ready removed in pre-processing (Section 2.3). It therefore captures votes that are broadly but not unanimously supported, rather than the chamber's truly consensual decisions, which are excluded by construction.

Table 6: Voting context summary. k^* is the number of voting contexts that best separates the votes, selected by the silhouette criterion over $k \in \{2, \dots, 10\}$. For EP6–EP7 it is a clear optimum; for EP8–EP10 it should be read as a lower bound, because the silhouette curve is nearly flat and its maximum moves to $k = 16$ –21 once the search is extended to $k = 40$ (Appendix C). Eff. k is the effective number of contexts, computed as the inverse Simpson index of the context-size distribution; unlike k^* it needs no cut-off. s is the silhouette score and p the null-model p -value.

EP	Votes	k^*	Eff. k	s	p
6	1,180	5	3.2	0.170	< 0.005
7	1,599	6	5.3	0.290	< 0.005
8	2,059	8	6.3	0.255	< 0.005
9	1,694	9	6.4	0.261	< 0.005
10	386	10	8.5	0.212	< 0.005

- 2. Pro-EU centre versus nationalist right** (96 votes, 25%). EPP + S&D + Renew + Greens + GUE versus PfiE + ECR + ESN. The “classic” Eurosceptic cleavage, now a minority pattern.
- 3. Left–right split** (75 votes, 19%). EPP + ECR + PfiE versus S&D + Greens + GUE, arising primarily on foreign-affairs votes. Includes a notable “horseshoe” sub-pattern in which PfiE and GUE/NGL jointly oppose the centrist mainstream on specific external-policy issues.
- 4. Green Deal fracture** (19 votes, 5%). EPP (+ 34 Renew) versus S&D + PfiE + ECR + Greens + GUE. EPP defects from the majority on Green Deal legislation, on which the rest of the chamber, including the Eurosceptic groups, remains aligned. This context has no precedent before EP9.

These four logics together account for 85% of EP10 votes. The remaining 15% are distributed across smaller contexts capturing trade consensus (GUE/NGL isolated), security and defence votes, and budgetary discharge clusters. The same party, notably the EPP, occupies different structural positions depending on the domain: anchoring the pro-EU majority on foreign affairs, voting with ECR on certain right-leaning resolutions, and defecting from the majority on Green Deal legislation. This behaviour is irreducible to any single ideological dimension.

5 Robustness

We subject the main conclusions to six robustness checks; Appendix E reports each in full. (i) Varying

the agreement encoding across binary, agreement-rate, cosine, linear, and alternative abstention weights leaves the directional trend intact and holds the EP10/EP6 modularity ratio between 3.7 and 4.4 (3.8 in the main specification), so the fourfold rise is not an artefact of the quadratic transform (Table 8, Appendix D). (ii) Varying the edge-weight threshold τ over $[0.05, 0.20]$ leaves the cross-term ranking of $Q_{\max}$ invariant. (iii) Bootstrap resampling of the per-topic analysis shows the Environment and Health estimate to be conservative rather than inflated. (iv) The multiplex context analysis survives a 200-permutation null model ($p < 0.005$ in every legislature), bootstrap resampling, and a label-noise perturbation test. (v) Down-sampling every legislature to EP10’s vote count reproduces the full-sample maxima within ± 0.003 and preserves the cross-term ordering, and normalising against a weight-preserving null model yields a size-corrected modularity $Q_{\max}/Q_{\text{null}}$ that still rises monotonically from 2.7 (EP6) to 6.4 (EP10). (vi) Restricting EP6–EP9 to their own first 19 months, matching EP10’s truncated record, leaves EP10 the outlier on entropy, opposition, and modularity alike; a vote-level bootstrap of its 386 non-consensus votes gives 95% confidence intervals that exclude every prior term ($Q_{\max} \in [0.088, 0.112]$).

6 Discussion

The analyses converge on a coherent narrative. The European Parliament has undergone a structural transformation from a consensus body to a high-entropy, multi-polar legislature. This transition did not occur discontinuously: three decades of European integration produced a parliament that could rely on a durable centre-leaning coalition, but accumulated shocks (the 2008 financial crisis, the migration crisis, Brexit, and the COVID-19 pandemic) progressively eroded the conditions sustaining that consensus. The chamber was already competitive along the left–right axis in EP6. Group cohesion around this period was high and rising [7, 8]. What faded was not competition itself but the broad, stable coalition that turned it into predictable majorities. What EP10 reveals is the consolidation of a new equilibrium.

This transformation runs along two distinct dimensions. The first is the erosion of consensus: winning coalitions are no longer broad, stable, or predictable, as the collapse in coalition-pattern concentration and the rise in opposition and vote entropy document. The second is polarisation: the chamber has hardened into well-separated blocs, with nearly half of all MEP pairs no longer sharing any stable alignment, visible in EP10’s bimodal agreement distribution. These are related but not equivalent, and the EP now occupies the low-consensus, high-

EP10 Voting Context Composition (NMI)




Figure 12: Party composition of voting blocs within each EP10 voting context (NMI). The ten silhouette-selected contexts are shown (a lower bound on the number of distinct geometries; see Appendix C). Four broad logics dominate: broad consensus (contexts 6, 7, 9), pro-EU versus nationalist right (contexts 4, 5), left-right split (contexts 1, 8), and Green Deal fracture (context 0).

polarisation regime rather than a polarised chamber that nonetheless retains a dependable governing majority. Importantly, this is not legislative gridlock: the Parliament continues to pass legislation, and a majority is still assembled on most votes. What has changed is the nature of that majority, which can no longer be assumed as a standing majority, but must be built, issue by issue, from shifting alignments. The *agreement crisis* we document is therefore a crisis of the old, stable-consensus model of majority formation, not of the Parliament’s capacity to legislate.

Our central empirical finding is that this transformation is both *systemic* and *domain-amplified*. It is systemic because every policy area for which we have sufficient data shows higher modularity in EP10 than in EP6, spanning the whole legislative agenda. It is domain-amplified because the Environment and Health area has undergone a disproportionate transformation, from structurally indistinguishable from the baseline in EP6 to more than double the parliament-wide average in EP10. The European Green Deal and the climate-policy agenda that preceded it are the strongest structural correlate of intra-coalition variance. Because the analysis is observational, this association cannot by itself separate the Green Deal from co-occurring developments (the von der Leyen investiture bargain, the Hungary–Poland rule-of-law disputes, and the COVID-19 recovery) or from the endogeneity of the agenda itself, since which environmental files reach a final vote is politically determined; we therefore read the Green Deal as having activated a latent

environmental cleavage that pre-dated it. Notably, this cleavage runs through the pro-European centre rather than along the centre–sovereigntist divide, consistent with the growing weight of the GAL–TAN axis (from green and libertarian to traditional and nationalist values) that underlies the transnational cleavage in European party competition [4, 34]: the Greens and much of the centre-left sit at its green pole, while the centre-right is pulled both ways.

This finding has implications beyond the EP. The Green Deal represents the most ambitious EU policy package since the single market, and its political effects appear to have begun restructuring the ideological space of the Parliament in ways unlikely to be reversed by adjusting individual provisions alone. The EPP’s fragmentation across three sub-communities, each separated along GMO policy, rule-of-law, and trade fault lines, reflects the heterogeneity of national parties within a formally unified group that the Green Deal has revealed and deepened. Renew Europe, by contrast, has converged with the EPP and S&D orbits: it remains internally cohesive (Section 4.2) but now votes with the grand-coalition centre, so community detection no longer resolves it as a separate voting force.

The growing Eurosceptic bloc exhibits a well-documented profile. It is, simultaneously, the most electorally significant development of EP10 and the least internally coherent main community: its cohesion ratio ($CR = 1.57$) is the lowest of all four main communities, and its voting behaviour is best characterised as coordinated opposition with no shared positive agenda, consistent with prior ac-

counts of Eurosceptic MEPs as a heterogeneous set that mostly opposes rather than builds a common programme [35, 36], spanning party families from the hard-Eurosceptic radical right to soft-Eurosceptic conservatives [37]. The multiplex analysis shows that this bloc participates in multiple coalition logics depending on the domain, joining the trade consensus, opposing rule-of-law mechanisms, and occasionally aligning with the radical left on specific foreign-policy votes. This heterogeneity limits its capacity to anchor an alternative majority but also gives the centre-right leverage: the EPP can selectively cooperate with parts of the right without formalising a structural alliance. Should this bloc continue to grow toward a numerical majority, the centre’s capacity to assemble winning coalitions on its own terms would itself become questionable – a contingency the remainder of the EP10 term will test.

A further question these dynamics raise concerns their downstream consequences. When majorities must be assembled anew for each vote from shifting cross-party alignments, the coordination required to secure them may increase the leverage of well-organised outside interests and of external actors seeking to influence specific votes. Whether this fragmentation has in fact heightened the Parliament’s exposure to lobbying or foreign influence is beyond the scope of a voting-structure analysis and is left to future work.

Several limitations qualify these findings. The EP10 data cover only the first 19 months of a five-year term. While all EP10 results pass statistical robustness tests, it remains possible that the legislature’s dynamics will shift as its vote record accumulates. The analysis is also restricted to final votes, which excludes the signalling and tactical dimensions of parliamentary behaviour expressed through amendments; final votes moreover mix legislative reports with non-legislative resolutions, which may behave differently (resolutions can isolate the far right more easily than technical legislation) and which we do not separate here. Finally, the per-topic analysis for EP10 is limited to six domains with sufficient votes; the remaining fourteen will become available for analysis as the term progresses. A further caveat concerns interpretation: rising modularity can reflect composition (more MEPs from structurally distant groups, especially after 2024), behaviour (groups voting less alike than before), or the agenda (more divisive files reaching a final vote), and the global $Q_{\max}$ series cannot separate these. The clearest evidence of behavioural change is the collapse of EPP–S&D agreement on Environment (Section 4.3), which occurs among groups present throughout the period and therefore cannot be compositional; in general, however, composition and behaviour remain entangled, since the sovereigntist groups of EP10 did not exist in EP6. Two further

caveats concern the policy-area analysis. First, the agenda is itself politically determined and its topic mix shifts across terms, so the composition of files reaching a final vote is not held constant; the systemic rise of modularity within every well-sampled policy area (Section 4.3), however, indicates that the hardening reflects within-domain change rather than the agenda drifting toward more divisive topics. Second, the policy-area labels follow the EP’s committee structure and are necessarily broad, aggregating heterogeneous files under single headings, so within-area heterogeneity is not resolved; nor were the labels themselves part of the cross-source validation of Section 2.1, which covered vote outcomes and member positions. Finally, we measure the structure of voting agreement, not legislative output: claims about the chamber’s capacity to govern are therefore indirect, and linking these structural shifts to legislative throughput and coalition size is a natural next step.

7 Conclusion

We have documented the structural evolution of the European Parliament over 22 years through a multi-method computational analysis. Four findings stand out.

First, the EP has experienced an *agreement crisis*: median opposition share more than doubled and coalition pattern predictability collapsed from EP6 to EP10, with no single voting configuration now accounting for more than 8% of final votes. Second, network modularity has increased fourfold, marking a genuine topological transition from a diffuse consensus network to a rigid bloc structure. Third, structural hardening is systemic across all policy areas but is disproportionately concentrated in Environment and Health, where the EU’s environmental agenda, from GMO authorisations to the European Green Deal, coincides with a modularity increase (0.012 to 0.211) that is unparalleled in magnitude or timing. Fourth, the EP’s coalition structure has fragmented from a few well-separated logics in EP6–EP7 into many weakly-separated ones by EP10, with the dominant logic now covering only 19% of votes. Together these indicate that EP10 is structurally distinct from the preceding two decades, pending confirmation as its vote record completes.

These results provide a quantitative baseline for monitoring the EP’s evolution through the remainder of the 10th legislature and beyond. The methodological pipeline, agreement scoring, network modularity, multi-resolution community detection, per-topic decomposition, and multiplex voting-context analysis, is fully reproducible and, with adaptation of the party-group and policy-area mappings, extensible to other national and supranational legislative assemblies exhibiting multi-party, multi-domain vote

records.

Declarations

Availability of Data and Code

All analysis code is openly available under the MIT licence at <https://github.com/ethz-coss/ep-voting-analysis>, including the scripts required to reproduce every figure in this paper. The repository bundles the harmonised roll-call matrices (`all_votes_main_EP6-EP10.csv`) and MEP membership records (`all_members_EP6-10.csv`) derived from VoteWatch Europe (now defunct; an archived copy of the site remains accessible through the Internet Archive) and the HowTheyVote.eu API. Users intending to redistribute the underlying vote data should consult the original providers' terms of use. All stochastic steps (Leiden restarts, bootstrap resampling, and permutation tests) use fixed random seeds, and exact package versions are pinned in the repository's `requirements.txt`, so every figure and statistic in this paper is reproducible from the released code.

Competing interests

The authors declare that they have no competing interests.

Funding

This work was partially funded with a grant from the Foundation for the University of Geneva.

Authors' contributions

NLR designed the study, assembled the dataset, performed the analyses, and wrote the manuscript. JS contributed to the study design and provided political science insights that shaped the interpretation of the results. DC and AR provided detailed feedback on the analyses and the manuscript. RB, DW, and DH supervised the research. All authors critically revised the manuscript and read and approved the final version.

Acknowledgements

The authors thank the maintainers of HowTheyVote.eu and the former VoteWatch Europe for making European Parliament roll-call records openly available.

References

- [1] European Union. Treaty on European Union (Maastricht Treaty). Official Journal C

191, 1992. URL <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:11992M/TXT>.

- [2] European Union. Treaty of Lisbon amending the Treaty on European Union and the Treaty establishing the European Community, signed at Lisbon, 13 December 2007, 2007. URL <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:12007L/TXT>. 2007/C 306/01.
- [3] Liesbet Hooghe and Gary Marks. A postfunctionalist theory of European integration: From permissive consensus to constraining dissensus. *British Journal of Political Science*, 39(1):1–23, 2009.
- [4] Liesbet Hooghe and Gary Marks. Cleavage Theory Meets Europe's Crises: Lipset, Rokkan, and the Transnational Cleavage. *Journal of European Public Policy*, 25(1):109–135, 2018.
- [5] Simon Hix, Abdul Noury, and Gérard Roland. *Democratic Politics in the European Parliament*. Cambridge University Press, Cambridge, 2007.
- [6] Richard Corbett, Francis Jacobs, and Darren Neville. *The European Parliament*. John Harper Publishing, London, 9th edition, 2016. ISBN 9780993454950.
- [7] Simon Hix, Abdul Noury, and Gérard Roland. Power to the Parties: Cohesion and Competition in the European Parliament, 1979–2001. *British Journal of Political Science*, 35(2):209–234, 2005.
- [8] Simon Hix, Abdul Noury, and Gérard Roland. Dimensions of Politics in the European Parliament. *American Journal of Political Science*, 50(2):494–511, 2006.
- [9] Andrew Scott Waugh, Liuyi Pei, James H Fowler, Peter J Mucha, and Mason A Porter. Party polarization in Congress: A network science approach. *arXiv preprint arXiv:0907.3509*, 2009.
- [10] Keith T. Poole and Howard Rosenthal. *Ideology and Congress*. Transaction Publishers, New Brunswick, NJ, 2007.
- [11] Fulvio Attinà. The Voting Behaviour of the European Parliament Members and the Problem of the Europarties. *European Journal of Political Research*, 18(5):557–579, 1990.
- [12] Thorsten Faas. To defect or not to defect? National, institutional and party group pressures on MEPs and their consequences for party group cohesion in the European Parliament.

- European Journal of Political Research*, 42(6): 841–866, 2003.
- [13] Darko Cherepnalkoski, Andreas Karpf, Igor Mozetič, and Miha Grčar. Cohesion and Coalition Formation in the European Parliament: Roll-Call Votes and Twitter Activities. *PLoS ONE*, 11(11):e0166586, 2016. doi: 10.1371/journal.pone.0166586.
 - [14] Gail McElroy and Kenneth Benoit. Party Groups and Policy Positions in the European Parliament. *Party Politics*, 13(1):5–28, 2007.
 - [15] Simon Hix. Legislative Behaviour and Party Competition in the European Parliament: An Application of NOMINATE to the EU. *Journal of Common Market Studies*, 39(4):663–688, 2001.
 - [16] Simon Hix and Bjørn Høyland. Empowerment of the European Parliament. *Annual Review of Political Science*, 16:171–189, 2013.
 - [17] Carlo Dal Maso, Gabriele Pompa, Michelangelo Puliga, Gianni Riotta, and Alessandro Chessa. Voting behavior, coalitions and government strength through a complex network analysis. *PLoS ONE*, 9(12):e116046, 2014. doi: 10.1371/journal.pone.0116046.
 - [18] Luc Stultiens. Community detection in the European Parliament: A network approach. Msc dissertation, University of Oxford, 2016. URL https://www.math.ucla.edu/~mason/research/Dissertation-Luc_Stultiens-final.pdf.
 - [19] Simon Hix, Abdul Noury, and Gerard Roland. Voting Patterns and Alliance Formation in the European Parliament. *Philosophical Transactions of the Royal Society B*, 364(1518):821–831, 2009.
 - [20] Nejat Arinik, Rosa Figueiredo, and Vincent Labatut. Multiple Partitioning of Multiplex Signed Networks: Application to European Parliament Votes. *Social Networks*, 60:83–102, 2020. doi: 10.1016/j.socnet.2019.02.001.
 - [21] VoteWatch Europe. European Parliament Voting Records 2004–2022. <https://www.votewatch.eu>, 2022. Dataset accessed via legacy archives; archived at <https://web.archive.org/web/20230107200751/https://votewatch.eu/>.
 - [22] HowTheyVote.eu. European Parliament Roll-Call Vote API. <https://howtheyvote.eu>, 2026. Real-time parliamentary data extraction.
 - [23] European Union. Consolidated version of the Treaty on the Functioning of the European Union, 2012. URL <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A12012E%2FTXT>.
 - [24] Clifford J Carrubba, Matthew Gabel, Lacey Murrah, Ryan Clough, Elizabeth Montgomery, and Rebecca Schambach. Off the record: Unrecorded legislative votes, selection bias and roll-call vote analysis. *British Journal of Political Science*, 36(4):691–704, 2006.
 - [25] Simon Hix, Abdul Noury, and Gerard Roland. Is there a selection bias in roll call votes? Evidence from the European Parliament. *Public Choice*, 176(1):211–228, 2018. doi: 10.1007/s11127-018-0529-1.
 - [26] European Parliament. Decision of 26 February 2014 on the amendment of Rule 166 of Parliament’s Rules of Procedure concerning the final vote and Rule 195(3) concerning voting in committee, 2014. P7_TA(2014)0146.
 - [27] Nikoleta Yordanova and Monika Mühlböck. Tracing the Selection Bias in Roll Call Votes: Party Group Cohesion in the European Parliament. *European Political Science Review*, 7(3): 373–399, 2015.
 - [28] Israel Mendonça, Rosa Figueiredo, Vincent Labatut, and Philippe Michelon. Relevance of Negative Links in Graph Partitioning: A Case Study Using Votes from the European Parliament. *The Second European Network Intelligence Conference (ENIC)*, pages 122–129, 2015.
 - [29] Vincent A. Traag, Ludo Waltman, and Nees Jan van Eck. From Louvain to Leiden: Guaranteeing Well-Connected Communities. *Scientific Reports*, 9(1):5233, 2019.
 - [30] Mark E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006. doi: 10.1073/pnas.0601602103.
 - [31] Catherine E. De Vries and Sara B. Hobolt. *Political Entrepreneurs: The Rise of Challenger Parties in Europe*. Princeton University Press, Princeton, NJ, 2020.
 - [32] Karl Rohe, Sourav Chatterjee, and Bin Yu. Spectral Clustering and the High-Dimensional Stochastic Blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
 - [33] Tai Qin and Karl Rohe. Regularized Spectral Clustering under the Degree-Corrected Stochastic Blockmodel. *Advances in Neural Information Processing Systems*, 26, 2013.

- [34] Liesbet Hooghe, Gary Marks, and Carole J. Wilson. Does Left/Right Structure Party Positions on European Integration? *Comparative Political Studies*, 35(8):965–989, 2002.
- [35] Nathalie Brack. *Opposing Europe in the European Parliament: Rebels and Radicals in the Chamber*. Palgrave Macmillan, London, 2018.
- [36] Duncan McDonnell and Annika Werner. *International Populism: The Radical Right in the European Parliament*. Hurst, London, 2019.
- [37] Aleks Szczerbiak and Paul Taggart, editors. *Opposing Europe? The Comparative Party Politics of Euroscepticism. Volume 1: Case Studies and Country Surveys*. Oxford University Press, Oxford, 2008. ISBN 9780199258307.
- [38] Alessia Amelio and Clara Pizzuti. Mining and Analyzing the Italian Parliament: Party Structure and Evolution. In *Recommendation and Search in Social Networks*, pages 249–279. Springer, 2015.
- [39] Arend Lijphart. The Analysis of Bloc Voting in the General Assembly: A Critique and a Proposal. *American Political Science Review*, 57(4):902–917, 1963. doi: 10.2307/1952608.

Appendices

A Resolution Selection for the Refined Community Analysis

The baseline community analysis (Section 4.2) uses the natural modularity resolution $\gamma = 1$, which yields a robust three-community partition in every legislature. The refined analysis instead adopts a higher resolution, selected per legislature by grid search to expose a party-group-scale (meso-scale) view. Figure 13 documents this choice. The Newman–Girvan modularity peaks at $\gamma = 1$ in all five terms (panel b), confirming the three-community structure as the natural optimum, while the number of recovered communities rises with γ (panel a). The refined resolutions used in the main text ($\gamma = 1.03$ – 1.13 , chosen to yield 8–10 communities matching the ~ 7 – 8 EP party groups; starred) lie on the rising flank of this curve rather than on a natural plateau, and are therefore best read as a deliberate meso-scale lens rather than an intrinsic community count. Beyond roughly 30 communities the partition fragments into singletons (over-resolution) and is not shown.

B Rank-Correlation Robustness of Cross-Topic Coherence

Figure 11 reports cross-topic coherence using Pearson correlations between per-topic agreement matrices. Because the Pearson coefficient assumes a linear relationship, Table 7 recomputes the mean coherence $\bar{r}$ per legislature with the rank-based Spearman and Kendall coefficients, with 95% confidence intervals obtained by bootstrapping over topic pairs (5,000 resamples). All three measures reproduce the same pattern, coherence peaking in EP7–EP8 and falling to its minimum in EP10, and the per-pair Pearson and Spearman values track each other closely ($\rho \approx 0.92$), confirming that the declining trend is not an artefact of the linearity assumption.

C Silhouette Stability of the Multiplex Context Count

Section 4.5 selects the number of voting contexts k^* by maximising the silhouette over $k \in \{2, \dots, 10\}$. To test whether the reported counts are artefacts of this search ceiling, Figure 14 extends the search

to $k = 40$. EP6 and EP7 show a genuine interior optimum ($k^* = 5$ and 6) that is unchanged by the wider search. For EP8–EP10 the silhouette is low and nearly flat, and its argmax drifts upward (to $k = 21, 16$, and 19 respectively) without a sharp peak, confirming that the main-text k^* for these terms are lower bounds rather than identified optima. The substantive conclusion does not rest on the precise count: the parameter-free measures of Table 6 (effective number of contexts; dominant-context share, which falls from $\approx 50\%$ in EP6 to 19% in EP10) establish the growing fragmentation independently of k .

D Encoding Robustness of the Modularity Trajectory

Section 5 reports that the fourfold rise in $Q_{\max}$ does not depend on the specific pairwise-agreement encoding. Table 8 makes this explicit, recomputing the full $Q_{\max}$ trajectory under the main quadratic encoding and six alternatives through the identical network-construction and Leiden pipeline ($\tau = 0.1$, $\gamma = 1$, best of five seeds). Every encoding preserves the same cross-term pattern, low and flat in EP6–EP7, a marked rise in EP8, a slight EP9 dip, and a peak in EP10, and every encoding yields an EP10/EP6 ratio between 3.7 and 4.4. The absolute $Q_{\max}$ level shifts with the encoding (the linear and cosine maps compress the score range and lower it uniformly), but the trend and its magnitude are invariant.

E Additional Robustness Checks

This appendix reports in full the checks summarised in Section 5. The encoding comparison is tabulated separately in Appendix D.

Agreement encoding. The agreement metric is varied across several alternative encodings drawn from the literature [17, 28, 38, 39], including binary (FOR/AGAINST only), standard agreement rate, cosine similarity, and alternative abstention weights; the full $Q_{\max}$ trajectory under each is reported in Table 8. All encodings reproduce the same directional trend: high agreement in EP6–7, a marked decline in EP8, a partial rebound in EP9, and a lower-agreement regime in EP10. Critically, recomputing the full $Q_{\max}$ trajectory under the linear encoding $A = (S + 1)/2$ (omitting the quadratic




Figure 13: Resolution profile of the global agreement network per legislature. (a) Number of detected communities versus the resolution parameter γ (shown up to 30 communities; beyond this the partition over-resolves into singletons). At $\gamma = 1$ every term yields the natural three-community structure; the shaded band marks the 8–10-community target and stars mark the refined γ used in the main text. (b) Newman–Girvan modularity Q versus γ , maximised near $\gamma = 1$.

Table 7: Mean cross-topic coherence $\bar{r}$ per legislature under Pearson, Spearman, and Kendall correlation, with 95% bootstrap confidence intervals (resampling over topic pairs).

EP	Pearson ($\bar{r}$ [95% CI])	Spearman ($\bar{r}$ [95% CI])	Kendall ($\bar{r}$ [95% CI])
6	0.656 [0.643, 0.668]	0.548 [0.534, 0.562]	0.416 [0.404, 0.427]
7	0.730 [0.712, 0.747]	0.665 [0.649, 0.681]	0.499 [0.485, 0.511]
8	0.727 [0.712, 0.744]	0.714 [0.696, 0.730]	0.536 [0.521, 0.551]
9	0.636 [0.617, 0.654]	0.640 [0.621, 0.657]	0.479 [0.464, 0.494]
10	0.493 [0.465, 0.520]	0.501 [0.472, 0.529]	0.374 [0.351, 0.396]

Table 8: Maximum modularity $Q_{\max}$ per legislature under alternative pairwise-agreement encodings. The quadratic row is the main-text encoding (Table 4); “binary” uses only FOR/AGAINST positions; “agreement rate” is the fraction of co-votes in agreement; the abstention-weight rows vary the ± 0.25 shared-abstention factor of Table 2. All encodings preserve the cross-term trend and a $\sim$ fourfold EP6 $\rightarrow$ EP10 rise. The EP10/EP6 ratios are computed from unrounded $Q_{\max}$ values.

Encoding	EP6	EP7	EP8	EP9	EP10	EP10/EP6
Quadratic (main)	0.026	0.033	0.054	0.053	0.100	3.8
Linear $A = (S+1)/2$	0.016	0.021	0.039	0.032	0.058	3.7
Binary (FOR/AGAINST)	0.025	0.034	0.059	0.056	0.110	4.4
Agreement rate	0.018	0.023	0.043	0.037	0.065	3.7
Cosine similarity	0.015	0.021	0.041	0.032	0.060	3.9
Abstention weight 0.5	0.027	0.033	0.056	0.055	0.103	3.8
Abstention weight 0	0.025	0.032	0.054	0.051	0.097	4.0

step of Eq. (2)) preserves the cross-term magnitude, not merely its direction: the EP10/EP6 ratio is essentially unchanged (3.7 under linear versus 3.8 under quadratic weighting), and every encoding yields a ratio between 3.7 and 4.4, confirming that the fourfold increase is not an artefact of the nonlinear transform.

Edge-weight threshold. The edge-weight threshold τ is varied over $[0.05, 0.20]$: the cross-term ranking of $Q_{\max}$ values is invariant.

Per-topic bootstrap. For the per-topic modularity analysis, bootstrap resampling (50 runs, 80% of votes) yields tight confidence intervals and confirms that the Environment and Health estimate is conservative at full sample size rather than inflated.

Multiplex context validation. The multiplex context analysis is validated against a 200-permutation null model ($p < 0.005$ in all legislatures), bootstrap resampling (mode k matches full-sample k^* within ± 1 in four of five cases), and a




Figure 14: Silhouette score as a function of the number of voting contexts k , per legislature, extended well past the $k = 10$ search ceiling (dotted line). Stars mark the published k^* . EP6–EP7 have stable interior optima; EP8–EP10 are flat plateaus whose argmax drifts upward, so their k^* are lower bounds.

5%-label-noise perturbation test for EP10 (mean ARI = 0.55, confirming that the gross structure survives noise).

Vote-count matching. To rule out that the near-fourfold growth in $Q_{\max}$ merely reflects EP10’s smaller vote record (386 non-consensus votes, versus 1,180–2,059 in earlier terms), we down-sample every legislature’s votes to EP10’s count while holding the MEP set fixed and recompute the modularity maximum over twelve independent subsamples. The matched-sample maxima reproduce the full-sample values within ± 0.003 for all five legislatures and preserve the complete cross-term ordering, with EP10 remaining roughly four times EP6 at an identical vote count; the $Q_{\max}$ trajectory is therefore a structural signal rather than a sampling artefact. Repeating this control on the full Leiden pipeline at native node count (667–714 MEPs) rather than a fixed-node proxy reproduces the published $Q_{\max}$ values to within ± 0.001 and yields the same conclusion.

Null-model normalisation. As a further safeguard against the size-comparability problem, we normalise each legislature’s $Q_{\max}$ against a weight-preserving null model: the agreement weights are randomly permuted and the network rebuilt under the identical τ threshold, so the null graph retains the same node count, edge count, and weight distribution

as the real one, but it carries no community structure; we recompute the Leiden modularity over 40 such nulls. The null modularity is uniformly small and nearly constant across terms ($Q_{\text{null}} = 0.010\text{--}0.016$), and, decisively, EP10’s null is not inflated by its smaller vote record ($Q_{\text{null}} = 0.016$, in line with the other legislatures). The size-normalised modularity $Q_{\max}/Q_{\text{null}}$ therefore rises monotonically from 2.7 (EP6) to 6.4 (EP10), more than doubling over the period, with every legislature exceeding its null by $z > 60$; the fourfold rise in $Q_{\max}$ is thus a genuine structural strengthening rather than a consequence of differing network size or density.

Early-term window. Because EP10 covers only the first 19 months of its term, we guard against an early-term *agenda* artefact (as opposed to the sample-size artefact addressed above) by restricting EP6–EP9 to their own first 19 months and recomputing the headline indicators. EP10 remains the outlier on every measure, highest vote entropy (57% versus 37–54% across the early-term slices), highest median opposition (0.25 versus 0.11–0.21), and highest modularity ($Q_{\max} = 0.100$ versus 0.028–0.071), even though early-term windows themselves tend to show *elevated* polarisation (EP6’s first-19-month $Q_{\max}$ of 0.052 already doubles its full-term value, reflecting contentious investiture and budget votes), which makes this a conservative test. EP10’s struc-

tural distinctiveness is therefore not an artefact of its truncated, early-term vote record.

EP10 vote-level bootstrap. A vote-level bootstrap of EP10’s own 386 non-consensus votes (300 resamples with replacement) yields tight 95% confidence intervals on its headline statistics that exclude every prior term: $Q_{\max} \in [0.088, 0.112]$, median opposition $\in [0.231, 0.269]$, and entropy $\in [56.2, 60.1]\%$.

F Voting Distance Between Political Groups

This appendix provides a complementary view of the fragmentation process at the level of group pairs. Figure 15 embeds the political groups of each legislature in a common plane based on pairwise voting disagreement: the central bloc that forms a single tight cluster in EP6 stretches into an elongated chain by EP10, while the EPP and the ECR move closer together.




Figure 15: Voting distance between political groups. For every divided main vote (winning side below 97.5% of the votes cast), each group is assigned its majority position (FOR, AGAINST or ABSTENTION; minimum five voters). The distance between two groups is the share of divided votes in which their majorities lie on opposite sides, counted with weight one half when exactly one of the two abstains or is tied. Each panel is a two-dimensional classical multidimensional-scaling embedding of these distances; panels are Procrustes-aligned across terms and share the same scale, so distances are comparable between legislatures. Circle area is proportional to seats at the constitutive session. Colours follow the group colour scheme used throughout the paper: EPP dark blue, PES/S&D medium blue, ALDE/Renew light blue, Greens/EFA green, GUE/NGL–The Left wine red, UEN/ECR orange, IND/DEM, EFD, EFDD and PfiE terracotta, ID and ESN dark brown. Non-attached members are excluded as politically heterogeneous. The grand-coalition partners drift apart (S&D–EPP distance 0.077 in EP6, 0.098 in EP8, 0.160 in EP10) while the EPP and ECR converge (0.415 in EP7, 0.290 in EP10).