

Two Crises, One Structure: Coupling and Scale in European Regional Ecosystems: supplementary material

Conway

Online supplementary material

Appendices A and B are intended for publication as online supplementary material.

Appendix A: Test ledger

This appendix lists every test run, what was asked, what was returned and the verdict carried into the paper. Grading is conservative, and borderline and null results are recorded as such.

A.1 Measurement and structure

Test	Question	Outcome	Verdict
PCA pillar composites (Stage 1)	Does each pillar form a coherent index?	PC1 captures 39 to 48 per cent within-pillar variance, PC1 vs equal-weight r 0.92 to 0.999	Pass, weighting-robust
One vs three factor (Stage 4)	Are the three pillars statistically separable?	Single general factor preferred by BIC by about 17 in both specifications (35 and 31 indicators), margin equals the parameter penalty	Three-factor not supported, conceded
CFA fit (Stage 4)	Does a clean three-factor model fit?	CFI 0.41 to 0.45, RMSEA 0.21 to 0.22, poor absolute fit for both structures	Documented limitation, with separability judged on the bifactor model and the typology agreement (Section 4.2)

Test	Question	Outcome	Verdict
One-factor shadow (Phase 1.3)	Is the typology just the general factor?	One-factor four-cluster vs published typology ARI 0.925, no single pillar reproduces it (ARI 0.39 to 0.54)	Typology one-dimensional but multi-domain-reliable
Triangulated anomalies	Do odd indicators recur across methods?	Overqualification and belonging proxy flagged in Stage 1, Stage 2, and Stage 4	Convergent, measurement doing real work

A.2 Typology

Test	Question	Outcome	Verdict
k selection (Stage 3)	How many tiers?	k = 4 chosen, silhouette 0.389, all four tiers stable	Four-tier ladder
Hennig Jaccard stability Centroid ordering	Are the tiers reproducible? Ladder or configurations?	0.83 to 0.91 in 2017 to 2019, all above 0.75 Centroids strictly ordered on all three pillars (-3.98, -1.57, +0.49, +2.17)	Robust Ordered ladder, configurational claim dropped
Calm-window persistence	Is the typology stable across two calm windows?	From 2014 to 2016 to 2017 to 2019, 86 per cent of regions stay in their tier, high-performing retention 0.94	Stable
Pre-pandemic replication (Phase: 2014 to 2016)	Is the typology a crisis artefact?	Four-tier structure present pre-pandemic, robust across four imputers and two seeds	Not an artefact
Spatial contiguity (Stage 5)	Are tiers geographically real?	Same-tier neighbour share 0.685 vs random near 0.28	Contiguous

A.3 The people-to-economy link

Test	Question	Outcome	Verdict
Annual coupling series (152-region panel, Section 4.5)	Where does the slide begin?	0.79 in 2008, 0.85 in 2011 to 2013, 0.80 in 2016, 0.76 in 2019, 0.73 in 2021 and 2023	The slide begins with the euro crisis recovery, not with the pandemic
Within-region imbalance (Section 4.5)	Do the bloc separation and the correlation agree?	Mediterranean economy minus people 0.39 in 2008, 0.03 in 2011, -0.24 in 2013, -0.30 in 2019. Spread of the gap across regions 0.64, 0.60, 0.73, then 0.82 in 2021, with the rise between countries (0.43 to 0.70) and none within (0.42 to 0.47)	The crisis first removed a boom-time excess of economy over people, then overshot, and the two measures agree from 2011
Shock overlap (Section 4.4)	Is the zero overlap arithmetic?	Worst quarter of both episodes: 0 of 152 on 2019 to 2023, 3 net of the 2019 displacement, 6 on 2019 to 2021, 9 within countries, against about 10 by chance	At or below chance in every version
Pre-pandemic calm-window baseline	Is high coupling the calm-time norm?	People to Economy about 0.79 to 0.83 across pre-pandemic windows, computed on 108 regions	Confirms calm-time coupling

The wide panel's before-and-after contrast in the same correlation, 0.847 to 0.776, is not used, for the measurement reasons given in Section 4.5.

A.4 Network (syndromes)

Test	Question	Outcome	Verdict
GGM estimation (Stage 2)	What is the partial-correlation structure?	EBIC-selected sparse network, 35 indicators, 85-edge backbone at the 0.12 threshold	Estimated
Community vs pillars	Are communities the pillars?	Four communities, ARI against the pillar partition at most 0.121	Cross-cutting syndromes

Test	Question	Outcome	Verdict
Syndrome strength	Stronger than chance?	About 5.5 times stronger than the upper bound of random partitions	Real structure

A.5 External validation and spatial

Test	Question	Outcome	Verdict
Income clubs (Phase 1.5)	Do tiers match income clubs?	Spearman 0.74 to 0.79, but tiers share income indicators	Sanity check, not independent
Development trap (Section 4.6)	Does the struggling tier coincide with the Diemer trap?	On the trap index built to specification on the annual panel, years trapped of fourteen are 7.2 struggling, 7.1 catching-up, 6.8 stable core, 6.9 high-performing	No relation between tiers and the trap, the two diagnoses are different conditions
Calm-window comparison (2014-16 to 2017-19)	Do the tiers hold their members between two calm windows?	Persistence 86 per cent, high-performing retention 0.94, upward moves outnumber downward about 2 to 1	Stable tiers, with net upward movement
Global Moran's I (Stage 5)	Are pillars spatially clustered?	0.84 to 0.90, all $p = 0.001$	Strong clustering
GNN cross-check (Stage 6)	Does graph structure add beyond tabular features?	Graph-blind R squared 0.88 to 0.91 vs GCN 0.82 to 0.83	Spatial signal already in features, no overclaim

A.6 Limits register

- The three-pillar structure is not statistically separable at the indicator level, so the framework's value lies in reliability and downstream outputs rather than in latent distinctness.
- The wide panel is a single cross-section and cannot identify dynamics. The annual panel describes them, and Section 4.8 identifies one cross-lag, from economy to people, within countries only and at a one-year horizon. Traps and the regulation that holds a region at its level are motivated here but not modelled, and they are left to the modelling companion.
- The development-trap index built to specification on the annual panel bears no relation to the tiers (Section 4.6). Roughly half of its nine comparisons fail by construction, so the absolute share trapped is not interpretable and only differences between groups and movements over time are.
- National context dominates the level of development (Appendices A.8 and A.9): between two

thirds and nine tenths of the gradient is between-country, depending on how finely the indicators are measured, and the typology is a cross-national ladder with within-country texture. The structure is a national-level phenomenon read at regional resolution.

- The within-country coupling differs markedly between countries (Appendix A.11), and with seven to forty-one regions per country the country-level figures are imprecise.
- The annual panel excludes the United Kingdom, Norway, Switzerland and Turkey, measures the people-side through education and health rather than trust and participation, and measures each region against the European mean of its own year, so a shock that hits every region equally is invisible to it by construction, and the panel records differentiation rather than damage. The pandemic has four years of aftermath against the crisis's eleven, so depths compare at matched horizons and recovery shares do not.
- The framework is structural and describes configurations. The political economy that produces and holds them, namely partisan governments, welfare regimes and redistributive choices, sits outside the object as a deliberate scope choice, distinct from the dynamic regulation left to the modelling companion.
- The institutional dimension is measured through a single government-quality index, which compresses welfare regimes and labour-market and financial institutions, so the place pillar is thin relative to the weight the argument places on institutions.
- The macro-financial drivers of divergence, foreign direct investment and core-periphery financial integration, operate over horizons a regional cross-section cannot capture and sit outside the design.
- The link from a region's position to political discontent is checked descriptively in Section 4.7 and Appendix A.10: anti-establishment vote shares fall monotonically up the typology, within countries as well as between them, with the far right concentrated in the catching-up tier rather than at the bottom, and years spent in the trap add nothing within countries. The check covers two elections and a short list of controls, under which the tier's association is shared with income per head, and the dynamic from trap to discontent stays with the literature (Rodríguez-Pose, 2018; Diemer et al., 2022) and with the modelling paper.

A.7 Integration with the pairwise literature (Section 4.9)

Test	Question	Outcome	Verdict
Marginal reproduction	Do the literature's cross-pillar couplings appear in these data?	All reproduce as raw correlations of 0.4 to 0.8 in the expected direction	Connected to the named findings
General-factor conditioning	Direct links or one shared dimension?	Of 224 strong cross-pillar correlations, the average falls from 0.50 to 0.18 on removing a single general development factor, 37 per cent drop to near zero	Most are the shared tide

Test	Question	Outcome	Verdict
Direct backbone	Is there structure beyond the tide?	66 direct cross-pillar edges, 56 cross-domain after removing 10 same-construct duplicates, 47 survive removing the general factor	A specific web remains
Caveats	What limits the claim?	Descriptive partial correlations, symmetric and non-causal, individual edges are small and fragile, so the claim rests on the aggregate, the same-construct split is a transparent judgment	Integrative, not causal

A.8 Within-country decomposition (the unit of analysis)

Test	Question	Outcome	Verdict
Variance decomposition	Is regional development national or regional?	Between-country differences account for 88 per cent of the general development factor on the pooled composites, within-country 12 per cent. On the same regions the share is 0.87 on the full set, 0.81 with the ten national-value indicators removed, and 0.66 on the regionally observed annual construct, the constructs agreeing on where regions stand at 0.93 (0.96 within countries)	National context sets the level, by between two thirds and nine tenths depending on measurement

Test	Question	Outcome	Verdict
Tier spread within countries	Are the tiers just countries relabelled?	68 per cent of the 25 multi-region countries span two or more tiers, several span three, average two tiers per country	The typology resolves within-country heterogeneity

A.9 Robustness to the indicator set

I specified three operationalisations in advance and report them together: the full set, the set with the most nationally-determined indicators removed (the survey-based social items and the government-quality index), and a register-only core. The checks are composite-level recomputations on observed data, and a full pipeline re-run per set is available but not required for these headline checks.

Operationalisation	Between-country share	Typology agreement (ARI vs headline)
Full set	83%	0.57
No survey or government-quality	79%	0.50
Register-only core	74%	0.41

The headline findings hold across all three: national context dominates the level, with the equal-weighted shares here (74 to 83 per cent) sitting inside the range of two thirds to nine tenths that the paper reports across all its constructs, and the typology is moderately stable. These are equal-weighted composites rather than the pooled-PCA composites of the main analysis, which is why the full-set between-country share here (83 per cent) corresponds to the 88 per cent of the pooled decomposition in A.8, and why the full-set typology agreement of 0.57 is a weighting-method baseline rather than a perfect match. The decline to 0.41 for the register-only core reflects a smaller indicator base rather than instability.

A.10 The typology and the geography of discontent

The table reports anti-establishment vote shares at the last national parliamentary election before the pandemic (2016 to 2019), from the European NUTS-level election dataset (Schraff, Vergioglou and Demirci, 2023), with parties classified as populist, far right or far left by the PopuList, version 4.0 (Rooduijn et al., 2024), matched on party identifiers, with 285 of 294 regions matched across 29 countries.

Tier	Regions	Anti-establishment share (mean, %)	Median	Far-right share (mean, %)
Struggling	39	40.8	42.5	9.3
Catching-up	63	40.3	41.6	25.3
Stable core	85	19.6	16.8	12.0

Tier	Regions	Anti-establishment share (mean, %)	Median	Far-right share (mean, %)
High-performing	98	15.6	16.8	11.6

- The gradient is monotone on the anti-establishment total. The Spearman rank correlation with tier rank is -0.46 raw ($p < 0.001$) and -0.27 with country means removed ($p < 0.001$), so the association is not only a country composition effect and the regional layer carries part of it.
- The far right alone is not monotone, since it peaks in the catching-up tier and is weakest in the struggling tier, whose anti-establishment vote flows more to the populist left. The two lower tiers hold similar volumes of discontent, but its political form differs.
- Against the trap index, on the 213 regions in 25 countries that carry both measures (Section 4.7), with the trap measured as the share of the years 2010 to 2019 spent trapped: within countries the anti-establishment share correlates with tier rank at -0.317 and with the trap at -0.009, and with country fixed effects the standardised coefficients are -0.295 (country-clustered standard error 0.074) on the tier and +0.014 (0.053) on the trap. The trap's raw association of +0.381 is entirely between countries. The far-right share shows little either way, -0.135 (0.199) on the tier and -0.008 (0.089) on the trap.
- At the previous national election (2011 to 2017), the mean anti-establishment share by tier is 36.3, 30.7, 11.0 and 14.0 per cent, with a Spearman correlation with tier rank of -0.39 raw and -0.19 with country means removed. With country fixed effects the coefficients are -0.119 (0.065) on the tier and +0.025 (0.034) on the trap. For the change in the share between the two elections they are -0.268 (0.102) on the tier and -0.013 (0.053) on the trap, and the mean change by tier is 5.0, 10.0, 9.8 and 1.6 points.
- With controls, on the 195 regions in 22 countries where density, net migration, income per head, unemployment and tertiary attainment are observed regionally (2017 to 2019 means), last election: the tier is -0.302 (0.080) with no controls, -0.261 (0.087) with log density and net migration, and -0.063 (0.059) once log income per head, unemployment and tertiary attainment are added, where income per head is -0.303 (0.107). The trap is +0.007, -0.007 and -0.027 across the three, never distinguishable from zero. All are produced by `code/analysis/vote_tiers.py`.
- One election further back (2006 to 2015), the mean share by tier is 31.7, 24.2, 11.2 and 16.5 per cent, with a Spearman correlation with tier rank of -0.20 raw and -0.01 with country means removed, and with country fixed effects the coefficients are 0.010 (0.083) on the tier and -0.019 (0.041) on the trap.
- With the fuller set of controls, on the 181 regions in 22 countries where all are observed, last election: the tier is -0.312 (0.083) with no controls, -0.261 (0.118) with log density, net migration, the share aged 65 or over, the share of foreign citizens and the manufacturing share of employment, and -0.143 (0.055) once log income per head (-0.205, 0.061), unemployment and tertiary attainment are added. Manufacturing is 0.115 (0.050) in the last of these. Without manufacturing the tier is -0.086 (0.062) on the same regions, so the tier's own association depends on holding manufacturing constant. The trap is 0.015, 0.018 and -0.016 across the three.
- With each measure taken before the vote (`code/analysis/vote_dated.py`), the level is the mean over the three years before the election of log income per head, or of an equal-weight composite of the fifteen annual indicators standardised within each year. The trap is the share of the years from 2010 to the year before the vote spent trapped, a history of six to nine years at the last election and one to seven at the previous one. Last election, 196 regions: income -0.257 (0.067) and trap +0.011 (0.047), composite -0.090 (0.042) and trap +0.034 (0.069), and

on 213 regions tier -0.295 (0.073) and trap +0.019 (0.055). Previous election: income -0.107 (0.047) and trap +0.043 (0.029), composite -0.068 (0.040) and trap +0.051 (0.034), tier -0.118 (0.062) and trap +0.036 (0.030). Third election, 207 regions, where the trap history is too short to enter: income 0.032 (0.071), composite -0.157 (0.053), and by pillar the people-side indicators -0.207 (0.069), against 0.027 (0.092) at the last election. The tiers are measured in 2017 to 2019 and so postdate the two earlier elections, which is why income and the composite before the vote are reported beside them.

- Caveats: vote shares are measured at three elections per country, the party classification is national, and the regressions are associations with a short list of controls. It corroborates the association of the lower tiers with anti-establishment voting (Dijkstra, Poelman and Rodríguez-Pose, 2020; Rodríguez-Pose, Dijkstra and Poelman, 2024) at the resolution of this paper’s typology. It does not test the dynamic from trap to discontent, and the negative result on the trap does not refute Rodríguez-Pose, Dijkstra and Poelman (2024), whose vote measure, trap horizon and specification all differ.

The natural extension is the 2024 European Parliament election. No NUTS 2 results file for it exists in the EU-NED release available (version 1, April 2022, national elections to 2020 and European elections to 2019) or in any official regional release I could cite, so the gradient is reported for 2016 to 2019 only and the extension is left open.

A.11 Differences between countries

This appendix gives a country-level view of the people-to-economy coupling, to show how far the pooled results of Sections 4.3, 4.5 and 5.1 vary from one country to another. It is descriptive, and with few regions per country the individual figures are imprecise.

The table gives the within-country correlation between the Human and Social Capital and Economic Activity composites in 2017 to 2019, for the sixteen countries with at least seven regions. Each figure is the correlation across that country’s regions. The pooled figure on the composites of the main analysis, 0.75, hides large differences in level. In Belgium, Czechia, Hungary, Norway, Poland and Romania the two sides move almost together within the country, with correlations above 0.9. In the United Kingdom the correlation is about 0.5, and in Greece and Austria it is close to zero or slightly negative. With between seven and forty-one regions per country the individual figures are imprecise, and those for Greece and Austria rest on thirteen and nine regions respectively.

Country	Regions	2017 to 2019
United Kingdom	41	0.54
Germany	38	0.68
France	22	0.77
Italy	21	0.90
Spain	18	0.75
Poland	17	0.91
Greece	13	-0.14
Netherlands	12	0.66
Belgium	11	0.95
Norway	11	0.98
Austria	9	-0.19
Czechia	8	0.94

Country	Regions	2017 to 2019
Hungary	8	0.92
Romania	8	0.96
Sweden	8	0.64
Switzerland	7	0.82

Appendix B: Figure and data provenance

- All wide-panel figures are generated from the validated analysis outputs under `data/clean/` in the analysis repository: pillar composites (`composites/pillar_scores_pooled.parquet`), typology (`clusters/cluster_consensus.parquet`, `clusters/cross_period_aligned.parquet`), structural model (`sem/composite_correlations_pooled.parquet`), spatial diagnostics (`spatial/morans_i.parquet`), and syndrome network (`networks/core_pcor_pooled.parquet`, `networks/core_communities.parquet`).
- Maps use Eurostat NUTS boundary files (1:1 million), 2016 vintage, projected to ETRS89-LAEA (EPSG:3035), joined on NUTS_ID. Tier colours are assigned by the strict pillar ordering of the centroids, so the legend reads low to high. Map join coverage: 287 of 294 regions, and the unmatched regions are vintage boundary changes rather than missing data.
- Tier labels (struggling, catching-up, stable core, high-performing) are assigned from the data by ranking the four cluster centroids on their mean pillar composite, not imposed.
- Figures from the wide panel: `fig_typology_map_core.png` (Section 4.3), `fig_network.png` (the partial-correlation connectogram, Section 4.2), `fig_decoupling_map_core.png` (each region’s standardised people-side score minus its standardised economy score in 2017 to 2019, Section 4.2), and `fig_tide_wiring.png` (the tide-versus-wiring slope chart, Section 4.9). The two maps, `fig_typology_map_core.png` and `fig_decoupling_map_core.png`, are built by `code/analysis/fig_paper2_core.py`, and the network and slope chart by `code/analysis/fig_paper2_figures.py`. Every figure is exported as a high-resolution PNG for the working draft and a vector PDF for final typesetting, all sharing one style file (`paper.mplstyle`) and one palette, Place `#2b6ca3`, People `#3a7d5c`, Economy `#a63d40`, with the tier ladder running from `#a63d40` through `#d08d5f` and `#8fb0c9` to `#2b6ca3`.
- The tide-versus-wiring figure is a slope chart over the 224 strong cross-pillar couplings (raw correlation at least 0.30). Each faint line runs from its raw correlation on the left to its strength after removing the single general development factor on the right, the bold line is the average, falling from 0.50 to 0.18. Lines plunging to the floor are absorbed into the shared dimension (37 per cent reach near zero), lines holding their height are the specific cross-domain web, and five exemplars are labelled, the institution-to-economy and trust-to-economy links collapsing, income-to-adult-learning, clean-air-to-social-contact and institutions-to-civic holding. Built from the same syndrome-network and standardised-panel outputs as the connectogram.
- The connectogram places indicators on a ring grouped by pillar, with edge colour for the sign of the partial correlation and bold versus faint edges for cross-pillar versus within-pillar links, so the cross-cutting syndromes are visible without a force-directed hairball.
- The country table in Appendix A.11 comes from `sem/within_country_by_country.parquet`, correlations across each country’s regions on the pooled composites of the main analysis.
- **The annual panel (Sections 4.4 to 4.8).** Fifteen indicators published by Eurostat at NUTS 2 in every year from 2008 to 2023, so that every value is a measurement of the region it is attached to and none is carried down from a national or regional aggregate. Eco-

nomic Activity: GDP per head in PPS (`nama_10r_2gdp`, logged), employment rate 20 to 64 (`lfst_r_lfe2emprrt`), activity rate (`lfst_r_lfp2actrrt`), unemployment rate (`lfst_r_lfu3rrt`), long-term unemployment share (`lfst_r_lfu2ltu`), and employment in high-technology sectors (`htec_emp_reg2`). Human and Social Capital: tertiary attainment and low education share, both ages 25 to 64 (`edat_lfse_04`), early leavers from education (`edat_lfse_16`), adult learning (`trng_lfse_04`), young people not in employment, education or training (`edat_lfse_22`), and infant mortality (`demo_r_minfind`). Place-based Conditions: life expectancy at birth (`demo_r_mlifexp`), road deaths per million inhabitants (`tran_r_acci`), and population density (`demo_r_d3dens`, logged). Indicators where a higher value is worse are sign-reversed before use.

- Regions enter if they are observed in at least 90 per cent of years on GDP per head, the employment rate and tertiary attainment. Gaps of up to three years are then interpolated and only regions complete on all fifteen indicators in all sixteen years are retained. This gives 152 regions in 23 countries, of whose cells 0.4 per cent are interpolated, concentrated in 2020 (1.8 per cent). Seven of the 152 have any gap in the 2008 to 2010 baseline. The United Kingdom is excluded because Eurostat ceased publishing its regional series, and Turkey because it is an outlier on several indicators. Norway and Switzerland do not meet the coverage rule. The balanced-panel rule drops regions whose codes changed during the window, which is why Poland contributes four of its seventeen regions. The blocs are Mediterranean (48 regions), Continental (61), Eastern (28) and Nordic (15).
- The composite is the first principal component of the fifteen standardised indicators in the 2008 cross-section, normalised to unit standard deviation there. Later years are scored on the same loadings and the same standardisation, so the scale is fixed while regions move on it. Each region is then expressed as a deviation from the European mean of its own year, which removes common growth and any shock that strikes all regions equally. Reported changes are therefore changes in relative position, in units of the 2008 cross-sectional standard deviation.
- **The development-trap index (Section 4.6)** follows Diemer, Iammarino, Rodríguez-Pose and Storper (2022). Three indicators, GDP per head in PPS (`nama_10r_2gdp`), nominal labour productivity per hour worked (`nama_10r_2nlp`) and employment over total population (`nama_10r_2emhrw` over `demo_r_pjanaggr3`), are each compared on five-year growth against three benchmarks: the population-weighted European aggregate, the region’s own country, and the region’s own preceding five years. DT1 is the share of the nine comparisons a region fails, and a region is counted as trapped when it exceeds one half. I chose the horizon, since the published source does not state theirs. This gives 236 regions in 27 countries for 2010 to 2023, including Serbia and four member states outside the 152-region panel. Because each comparison is made against an average or a previous period, roughly half fail by construction and the median region sits near the threshold, so differences between groups and movements over time are interpretable and the absolute level is not.
- Both are built by `code/processing/build_annual_panel.py`, which pulls from the Eurostat dissemination API with a sidecar recording table, filters and pull date for each series, and writes `data/clean/annual_panel/panel_final_2008_2023.parquet` and `data/clean/annual_panel/dt1_k5.csv`. The annual-panel figures, `fig_two_shocks_blocs.png` and `fig_two_shocks_pillars.png` (Section 4.4), `fig_coupling_annual.png` (Section 4.5) and `fig_trap_bloc_tier.png` (Section 4.6), are built by `code/analysis/fig_two_shocks.py`.
- **The cross-lagged regressions (Section 4.8)** are built by `code/analysis/cross_lags_annual.py` on the same 152-region panel, with output in `output/tables/cross_lags_annual.csv`. Coefficients are in pooled standard deviations of each composite, with each region’s own level and each country’s common movement removed, and country-clustered bootstrap intervals from

500 draws.