A Shock Reveals the Structure: Coupling, Typology, and Development Paths across European Regional Ecosystems

Conway

Contents

Extended summary

This paper tests, on European regions at NUTS 2 level, the ecosystem framework set out in the conceptual companion. That framework organises regional development around three interacting pillars, **Place-based Conditions**, **Human and Social Capital**, and **Economic Activity**, and proposes that a region is a coupled, dynamic system rather than a stack of separate factors. The test produces two main findings. First, most of the variation in development across European regions is in fact variation between countries: between-country differences account for roughly three-quarters to nine-tenths of the total, and under the pandemic the correlation between the social and the economic sides of development fell between national systems while within countries it held steady. Second, the four tiers of regions that emerge from the data hold on to their members very unequally: the struggling tier kept more than nine in ten of its regions across the shock, while the high-performing tier lost nearly half, so the regional structure is far stickier at the bottom than at the top.

The overarching idea is that a tightly coupled system looks one-dimensional in calm times and shows its separate parts only under stress. In normal conditions the three pillars do not behave as three separable dimensions: a single general factor fits the thirty-five indicators better than three factors do, and the pillar composites correlate at 0.85 to 0.89, high enough that, statistically, they could be read as one underlying development 'factor'. When the pandemic hits, however, the correlation between Human and Social Capital and Economic Activity falls while the two correlations involving Place-based Conditions barely move. Dimensions that were truly independent would not have correlated so highly in calm times, and indicators that were really just one factor could not have come apart at one specific link under stress, so the two results together point to a coupled system. This is why

the single-factor result, which taken alone could be read as a refutation of the framework, is in fact its first piece of supporting evidence. You cannot see the structure of a regional economy until the system is stressed.

The four tiers themselves (struggling, catching-up, stable core, and high-performing) are stable across the pandemic shock, geographically contiguous, and already present in the pre-pandemic window, so the typology is not an artefact of the crisis. Their unequal grip on their members is what connects them to the regional economics of clubs and traps: the tiers behave like basins of differing depth, and movement between them across the shock ran downhill, with downward transitions outnumbering upward ones by more than four to one. The pattern resembles allostasis, in which a region's configuration is the most stable setting its conditions allow rather than a good one. This is why a struggling region can be stable and poor at the same time, and why effort alone does not move it. The tiers also map onto the geography of discontent: in the last national elections before the pandemic, anti-establishment parties took about twice the vote share in the two lower tiers that they took in the two upper ones, a gradient that holds within countries as well as between them (Appendix A.10).

The national share of the variation is no surprise once one considers that institutions, welfare regimes, language and politics are largely national, and it survives dropping the most nationally-determined indicators, so it is not an artefact of measuring national things. It does change what the typology is: the four tiers form a cross-national development ladder read at regional resolution, with the regional layer adding texture within countries rather than driving the differences. It also locates the decoupling: the fall in the people-to-economy correlation under the pandemic happened between national systems, while within countries the correlation held steady or tightened in every indicator set tried. What the pandemic pulled apart was the alignment between national systems, which points to their divergent policy responses.

A final result connects back to the conceptual companion. The pairwise findings the field has accumulated, government quality and trust, social capital and skills, institutions and growth, all appear in this panel, and most of them turn out to reflect a single shared development dimension: of 224 strong cross-pillar correlations, more than a third drop to near zero once that shared dimension is removed, while 47 direct links survive it. The method behind this decomposition, a partial-correlation network estimated over the development indicators themselves, is borrowed from psychometrics and is, to my knowledge, new to regional studies. The framework does not predict these findings, but it shows that they belong to one structure.

The paper is descriptive, in the tradition of the clubs and traps work it builds on. It measures what the structure of European regional development is and how it moved under a common shock, identifying what causes what, and measuring how people feel, both lie beyond its design. Three snapshots can describe these paths but cannot model them. The fall in the people-to-economy correlation is moderate, 0.85 to 0.78, and borderline in significance on any single test, the conclusion rests on four lines of evidence agreeing in direction rather than on one decisive number, and since the comparison spans roughly twenty-five national systems over a single shock, it is an association, not a causal claim. Because so much of the variation is the European east-west income gap, the national finding describes economic geography at least as much as institutions. The mechanism that holds a region in its basin is left to the modelling paper that follows.

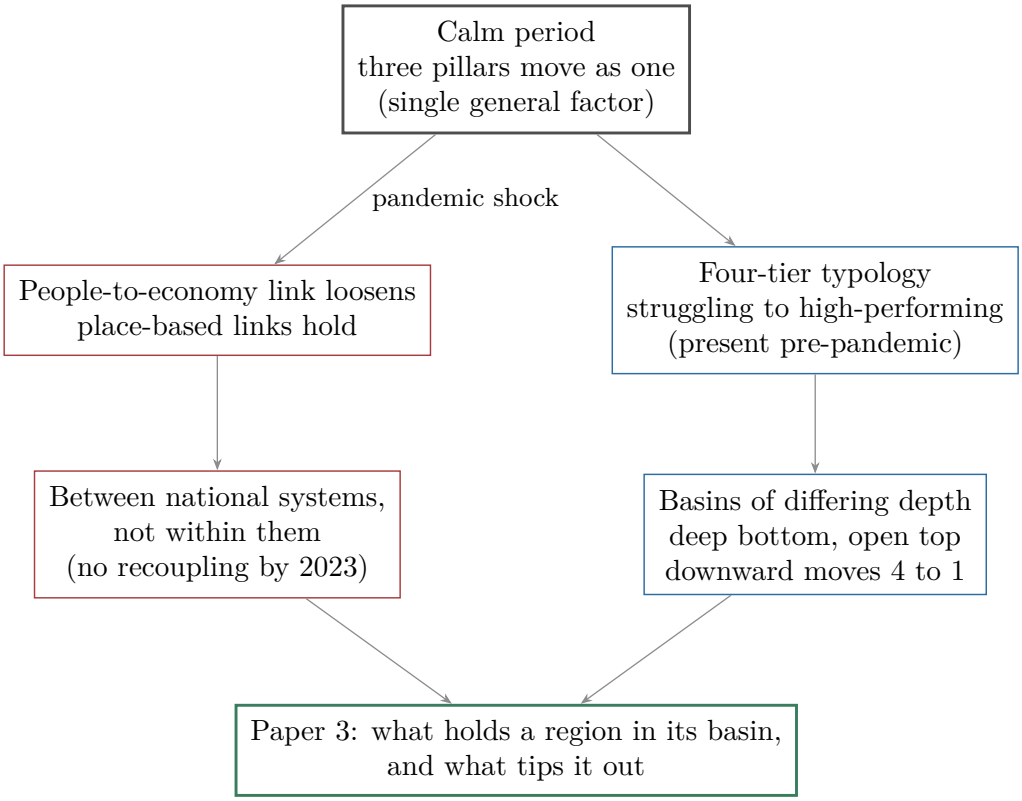

Figure 0. The argument in one view: the calm-time single factor and the selective decoupling under shock are two faces of one coupled system, the typology beneath behaves as basins of differing depth, both findings hand their starting conditions to the modelling paper.

Section 1: The question, the idea, and the design

1.1 What the paper asks

The conceptual companion sets out the ecosystem framework as a proposal, a way of seeing regional development as a coupled system rather than a stack of separate parts, but it does not test it. That is the task of this paper, which puts a direct question to the European regional data: does the framework hold, and if it does, what does a region as a coupled system actually look like? Answering this requires being clear from the outset about what is being measured and at what scale.

- The lens throughout is wellbeing in the regional economy. We look past output to the people-side of a region and the territorial conditions beneath it, the things that aggregate figures miss, following the decoupling of wellbeing from wealth that motivates the companion (Lima de Miranda and Snower, 2020; Conway, 2020). The people-side is measured through human and social capital, education, skills, labour-market inclusion, institutional trust, civic participation and belonging. These are treated as a window onto the social side of development, not as a direct measure of how people feel, so the loosening documented later is a loosening between measured conditions rather than a reading of subjective wellbeing.
- The unit is the European region at NUTS 2 level. This is the scale at which Cohesion Policy actually operates, and the scale at which all three pillars can be observed at once.

- The framework is taken at its word, as a set of testable commitments rather than a metaphor: three organising domains, named mechanisms running between them, and the claim that the system is coupled and settles into a small number of configurations that are stable but able to shift.

What follows allows the data to qualify the framework as much as to confirm it. The most important qualification, that the three pillars behave as a useful organising structure rather than as three cleanly separable dimensions, runs through everything that follows.

1.2 The organising idea

The core of the paper is a pair of results that pull in opposite directions until they are read together. Because the whole argument depends on holding them together, I set them out here before the evidence.

- In calm times the three pillars behave almost as one: across the 294 regions in 2017 to 2019, a single general factor fits the indicators better than three separate factors do.
- Under the pandemic, the correlation between the people-side and the economy falls from 0.85 to 0.78, while the two correlations involving place decline by only about 0.03 each.
- Independent dimensions would not have produced the first result, and a single underlying dimension could not have produced the second, so together the two results indicate a coupled system.

This converts the single-factor result from an apparent refutation of the framework into supporting evidence, since strong calm-time correlation is exactly what a coupled system should produce.

1.3 Data and design

The evidence rests on a harmonised panel of European regions assembled for this purpose, wide enough to see all three pillars at once and reaching across the shock.

- The panel covers 294 NUTS 2 regions and thirty-five indicators spread across the three pillars, over three windows: a core period (2017 to 2019), a post-shock period (2021 to 2023), and a pre-pandemic period (2014 to 2016) kept for validation. The indicators come from Eurostat regional statistics, the European Social Survey, the European Quality of Government Index, and satellite-derived PM2.5, harmonised across the 2013, 2016 and 2021 NUTS vintages.
- The indicator set is theory-led. Each indicator is there to measure one of the three pillars, drawn from the standard regional sources, and the set is then trimmed for coverage, dropping anything too sparse to impute, and for redundancy. There is also a plain practical constraint: the multivariate methods need a single fixed, common, complete-enough matrix, and cannot take a different set of indicators region by region, so a set of roughly this size is dictated as much by the method as by choice. It is one of several reasonable ways to operationalise the pillars, and Appendix A.9 shows the main findings survive three of them, the full set, the set stripped of its most nationally-determined indicators, and a register-only core.
- Missing values are filled by random-forest multiple imputation (MICE, twenty draws) and pooled by Rubin's rules, 8.5 per cent of cells in the core period and 18.8 per cent post-shock are imputed, the latter figure inflated by the thinning of UK data after Brexit.
- The analysis then runs as a staged pipeline, and each stage is a separate test of the framework: principal-component pillar composites, a Gaussian graphical model of the indicator network, a

consensus clustering with stability testing, a composite-level structural model, spatial diagnostics, and a causal forest. The reason for using so many methods is that the framework is then asked the same question in several independent ways, and the answer can rest on whether they agree rather than on any one of them.

- One methodological choice is worth flagging, because it is new to this field. Building a partial-correlation network of development indicators at NUTS 2 does not seem to have been done before in regional studies. The method is well established in psychometrics (Epskamp and Fried, 2018; Haslbeck and Waldorp, 2020) and in systems biology, and the graphical lasso has been used in economics on networks of places, such as inter-country financial risk, but there the places are the nodes. Here the nodes are the development indicators themselves, which is what lets the syndromes be read straight off the network.

The design is deliberately redundant in this way, so that no single method is left to carry the conclusion alone.

Section 2: Strong coupling in calm times

2.1 Regional development is strongly and spatially structured

Before considering whether this structure constitutes an ecosystem, it is necessary to establish that there is a strong, organised structure to account for in the first place. The data show that there is, and that it is pronounced.

- Each pillar holds together as a coherent index. Its first principal component captures between 39 and 48 per cent of the variation within the pillar, and these PCA-weighted composites correlate between 0.92 and 0.999 with a simple average of the same indicators, so nothing that follows depends on how the indicators are weighted.
- The three pillars move closely together across regions. In the core period place and people correlate at 0.88, place and economy at 0.89, and people and economy at 0.85.
- Every pillar is strongly clustered in space, with a global Moran's I between 0.84 and 0.90 in the core period and between 0.80 and 0.86 afterwards, all comfortably significant, and the typology that emerges is geographically contiguous as well, neighbouring regions falling in the same tier about seven times in ten against a chance rate closer to three in ten.

Regional development is therefore strongly and spatially organised, which is not in dispute. On its own, however, this establishes little. Coherent composites, close correlations and spatial clustering are all equally consistent with a single development gradient running across Europe, and the calm-time picture does not reveal whether that gradient is best understood as one dimension or as a coupled system of three. That question cannot be resolved here, it is answered by the system's behaviour under stress, which is the subject of Section 4.

A reliable composite rather than three separable dimensions

The framework asserts that its three pillars constitute three distinct dimensions of regional development. The data qualify that assertion, and in an instructive way.

- At the level of individual indicators, the three pillars simply cannot be told apart. A single general factor is preferred over a three-factor structure by a BIC margin of about 17 in every specification, and since that margin is almost exactly the penalty the criterion charges for the

extra parameters, the three-factor model buys no real improvement in fit. This is a comparison of relative parsimony, not of absolute fit. Both structures fit the indicator data loosely, as the confirmatory model in Appendix A.1 records, so the point is narrow: splitting the pillars into three factors adds parameters and returns nothing for them.

- What structure the indicators do contain runs across the pillars, not along them. The Gaussian graphical model recovers communities of indicators, cross-cutting syndromes, that bear almost no relation to the pillar boundaries, their agreement with the pillar partition never rises above an adjusted Rand index of 0.12. Indicators group by what they do, in other words, not by which pillar they were assigned to.

The syndrome network: partial correlations at |r| ≥ 0.12,
bold edges cross pillar boundaries

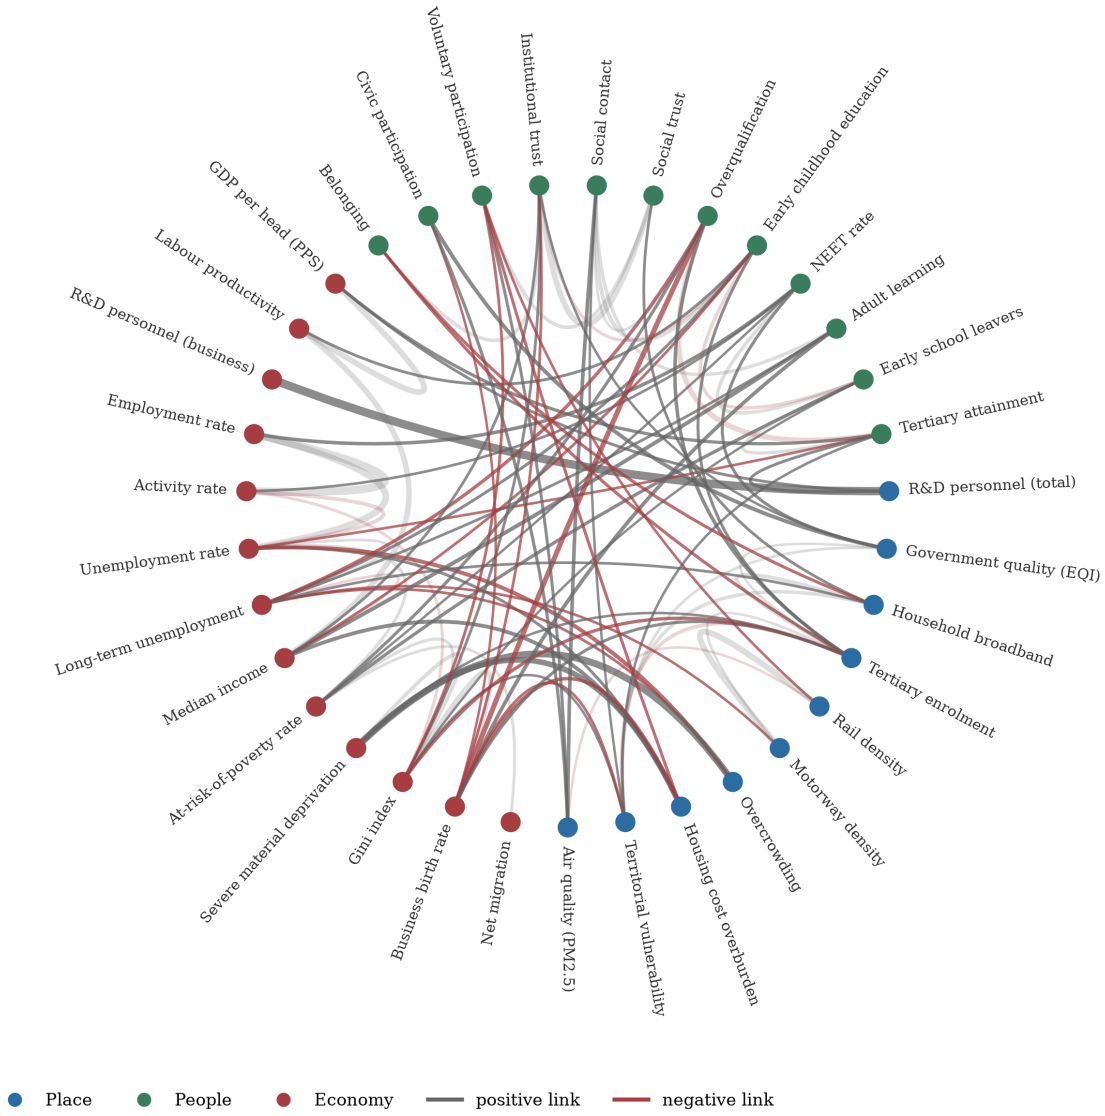

- Nevertheless, the three-pillar composite achieves what no single pillar can. The typology it produces lies along a single development factor, with which it agrees at 0.89 to 0.93, yet no individual pillar reproduces that typology on its own, agreeing with it only at 0.39 to 0.54. Only

the three pillars combined recover it, because averaging across them cancels the idiosyncratic noise present in each. The composite is, in short, more reliable than any of its components, even though those components are not statistically separable.

- The same two oddities turn up wherever one looks, which is usually a sign that a measurement is picking up a real signal: overqualification and the belonging proxy load against the people-side composite, reappear as anomalies in the network, and stay anomalous in the confirmatory model.
- A single factor does not, however, account for everything. The map below presents, for each region, the difference between its people-side standing and its economy-side standing in the post-shock period. This difference has a clear and stable geography, with the people-side ahead across the Nordic and Iberian periphery and the economy ahead across the German and North-Italian core. Whether it reflects a real distinction between people and economy, or merely a single factor with uneven loadings, cannot be determined from the map, that question is settled in Section 4. This is the configuration on which the shock then acts.

Where the people-side runs ahead of the economy, post-shock

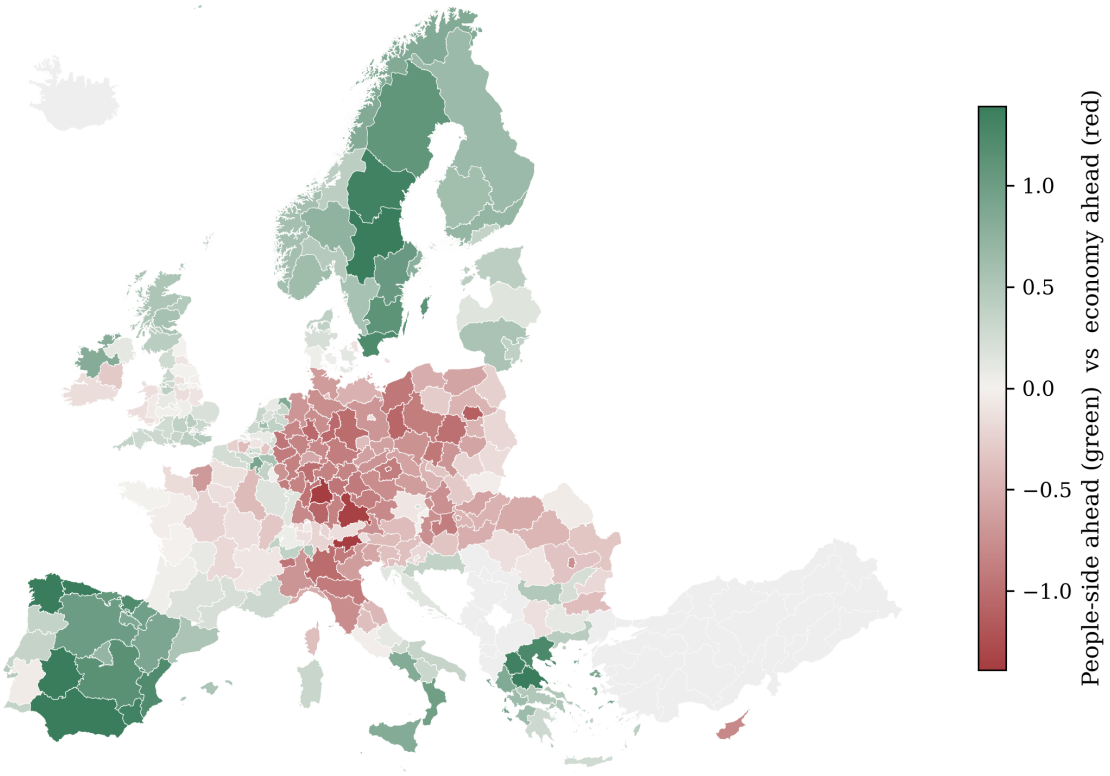

The pillars are therefore best understood as a conceptual organisation and a reliable composite rather than as three separable dimensions of regional development. Their value lies less in statistical separability than in what they produce, namely a stable typology, a layer of syndromes, and a selective response to the shock. This is a real limitation of the framework, and the rest of the paper builds on it.

Section 3: The typology and the syndromes

The most concrete product of the framework is a four-tier typology of regions, which proves to be stable, ordered, and geographically coherent.

- The four tiers, struggling, catching-up, stable core and high-performing, contain 39, 65, 92 and 98 regions respectively in the core period. Their centroids are strictly ordered on each of the three pillars taken separately: struggling at Place -3.7, People -3.6 and Economy -4.6, catching-up at -1.5, -1.8 and -1.4, stable core at +0.5, +0.4 and +0.5, and high-performing at +2.0, +2.3 and +2.3. It is this ordering on every pillar, rather than any single aggregate score, that warrants reading the typology as a ladder.
- The ladder is stable across the shock. Each tier attains a Hennig Jaccard bootstrap stability above 0.75 in both periods, between 0.83 and 0.91 in the core period and between 0.86 and 0.99 afterwards, and 79 per cent of regions remain in their tier from one period to the next.
- Its geography is the one a regional economist would anticipate: the high-performing regions lie in the north and west, the struggling regions in the south and east, and the stable core visibly expands after the shock.

The four-tier typology, before and after the shock

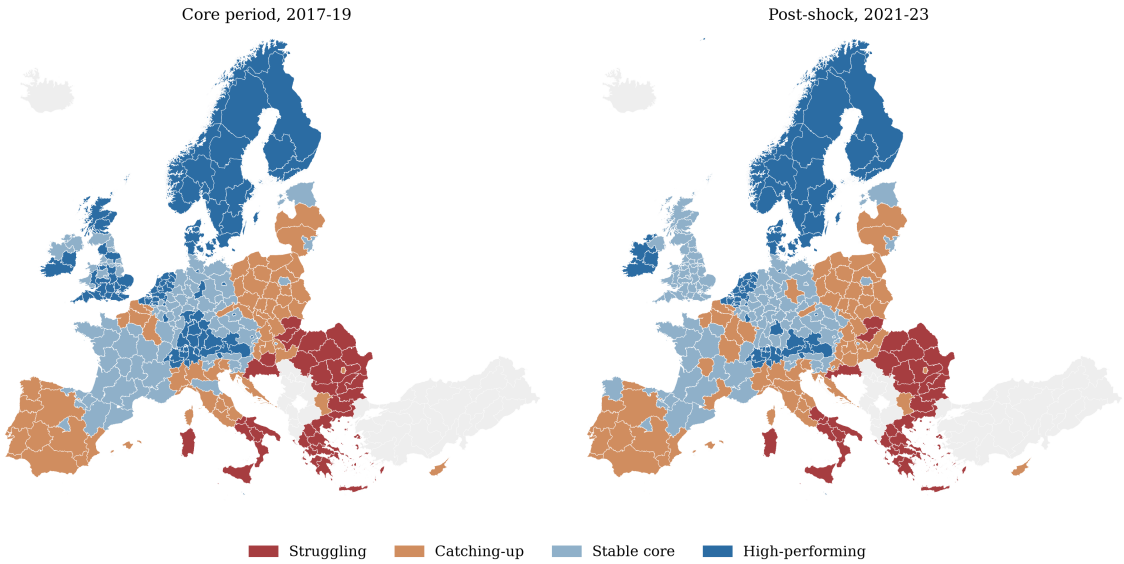

- Beneath the ladder lies a layer that a single composite index would not reveal: cross-pillar syndromes, four in the core period and six afterwards, each between 5.5 and 8.3 times stronger than a random grouping of the indicators. This descriptive structure is available only because the measurement spans three domains.
- The typology is, more precisely, a cross-national development ladder observed at regional resolution. National context dominates the level: between-country differences account for the large majority of the development gradient, between three-quarters and nine-tenths depending on the measure (74 to 88 per cent across the pooled decomposition and the robustness operationalisations of Appendices A.8 and A.9), which is the expected consequence of national institutions in a panel spanning the European East-West gap. The regional unit nonetheless contributes real texture, since two thirds of the multi-region countries contain regions in more than one tier, several in three. And the decoupling examined in Section 4 resides in this between-country

component: across national systems the people-to-economy alignment loosened, while within countries it held, so that the loosening under the shock is a divergence between national systems rather than a within-country regional dynamic.

The typology is thus robust, ordered and geographically coherent. Section 5 follows the regions as they move between its tiers.

Section 4: The decoupling under the pandemic

The pandemic did not move all parts of the regional economy together: one link weakened while the others held, a difference that any single aggregate measure would miss. No single piece of evidence establishes this on its own. Rather, four lines of evidence, drawn from four different methods applied to the same panel, point in the same direction. They are not statistically independent, since they share the indicator base, the imputation procedure and the period contrast, but each contributes something the others cannot: the composite correlation supplies the structural estimate, the tier-specific pattern shows where the loosening concentrates, the typology shift records its distributional consequence, and the changes in the network indicate which indicator pathways transmitted the shock. It is their convergence, rather than any one of them alone, that carries the conclusion.

- The main result is a selective loosening. Between the core and post-shock periods, the correlation between the people-side and the economy falls from 0.85 to 0.78, while the two correlations involving the place-based foundation barely move, each declining by about 0.03. The people-to-economy link loosens more than twice as much as the others, which is precisely the relation the framework predicts should give way (Wald z = 2.51, p = 0.012).

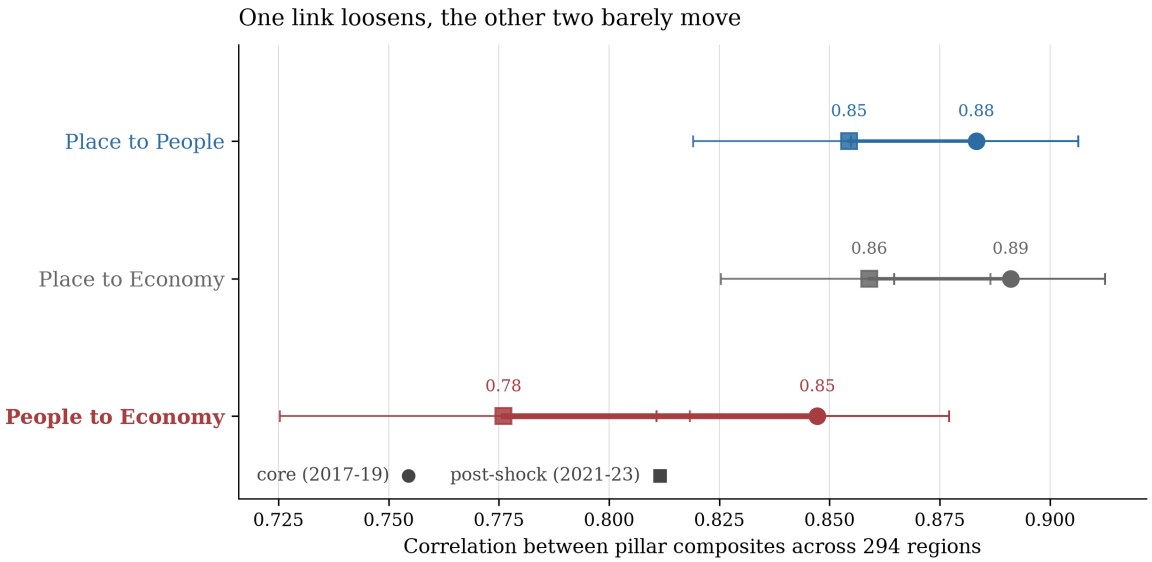

The people-to-economy coupling, before and after the shock

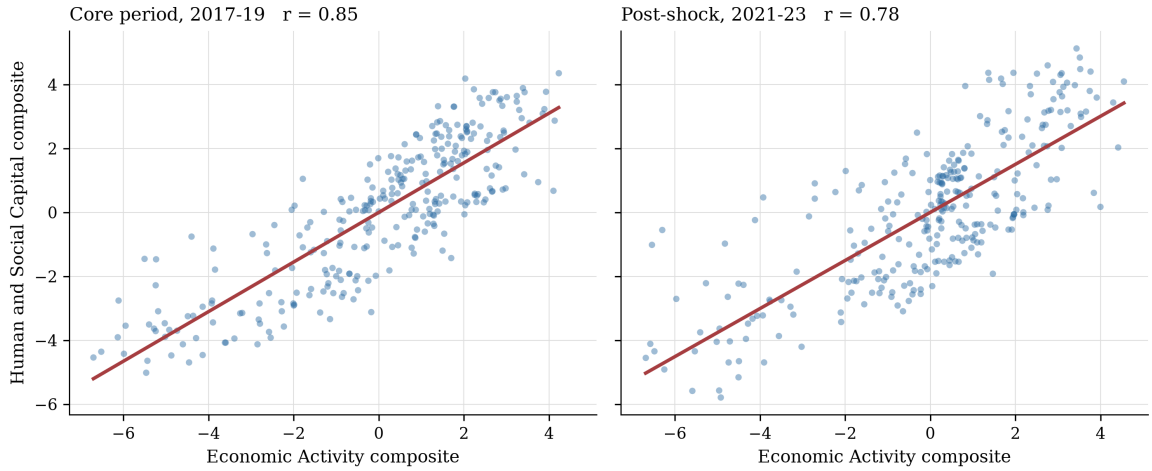

The structure behind this selectivity matters, since the people-side enters two of the three pillar relations, and any uniform weakening of a common factor would therefore loosen its link to the economy and its link to place together. The data instead show a change confined to the people-to-economy relation while the people-to-place relation holds, which no single-factor account with stable loadings can produce, though a coupled system can. A one-factor model whose loadings shift with the shock could mimic it, but at that point it has stopped being the one-factor story.

The natural objection is that this is not a real loosening but an artefact of the way the pandemic disturbed different economic indicators. Three forms of this objection can be examined, and none of them holds.

1. The first concerns furlough. Job-retention schemes sustained measured employment and held unemployment down, so the decoupling might reside entirely in those distorted labour-market figures. This turns out not to be the case: when the four labour-flow indicators are removed altogether, the link still falls, from 0.79 to 0.74, and when the economy is narrowed to its hardest core, output, productivity and household income, the quantities that actually contracted and that furlough cannot inflate, the loosening is sharper still, from 0.72 to 0.62. If furlough were responsible, removing the labour flows would dissolve the result, instead it is the most purely economic measures that display it most strongly.
2. The second concerns the United Kingdom, whose data thinned after Brexit. The result survives the exclusion of the UK, holds in every leave-one-country-out test, and scarcely changes under country-clustered inference.
3. The third concerns the imputation used to fill the missing data. The decline holds on a fixed set of 124 regions with no missing values at all, from 0.79 to 0.59, and the furlough re-analysis above uses observed data only. That complete-case set covers a narrower geography, which limits how far it generalises.

A fourth question concerns the scale at which the loosening occurs, and the answer shapes actually how the result should be read. The loosening is a between-country divergence, as we see with each region's national mean removed, the people-to-economy correlation within countries is flat to slightly rising across the shock, between about 0.65 and 0.76 in every indicator set, while the raw decline resides entirely in the between-country component. The pandemic therefore loosened the people-to-economy alignment between national systems while leaving it intact within them, which

points to divergent national responses to a common shock rather than to a within-country regional dynamic. The regional resolution still matters, since the standing structure and the typology carry real within-country detail, but the loosening under the shock happened at the national level.

How much weight can a single result of this kind bear? Considered in isolation it is borderline by the strictest standard, with a clustered Bonferroni p between 0.05 and 0.06, and the argument does not rest on that threshold. Two considerations matter more. First, the framework identified the people-to-economy relation in advance as the one that should give way, so this is a test of a stated prediction rather than a search across many pairs. Second, the correlation is stable in calm times and moves only under the shock: on a consistent measure it stands at about 0.8 in both the pre-pandemic and the core periods, two readings that scarcely differ, and falls into the mid-0.6s only afterwards. A single p-value does not capture a before-and-after pattern of this kind, the pattern is what distinguishes a structural loosening from noise.

The window contrast can also be taken year by year. The figure below plots the people-to-economy correlation annually from 2014 to 2023, on a fixed balanced panel of 139 regions in 25 countries, using the register-only core so that every year is measured the same way. The raw series drifts gently down through the calm years, from 0.80 in 2014 to about 0.76 in 2019, then steps down through the pandemic to 0.69 in 2021, the two pandemic years removing more than the previous five combined. Nor has it returned: in 2023 the correlation still stands at 0.69, three years after the shock. The within-country series holds between 0.76 and 0.81 across the whole decade, with no pandemic step at all. The confidence bands are country-clustered, wide, and overlap from year to year, so the trace describes a trajectory rather than testing one. But the trajectory says what the windows said, and it adds one thing the windows could not: the loosening between national systems has persisted to the end of the observation window.

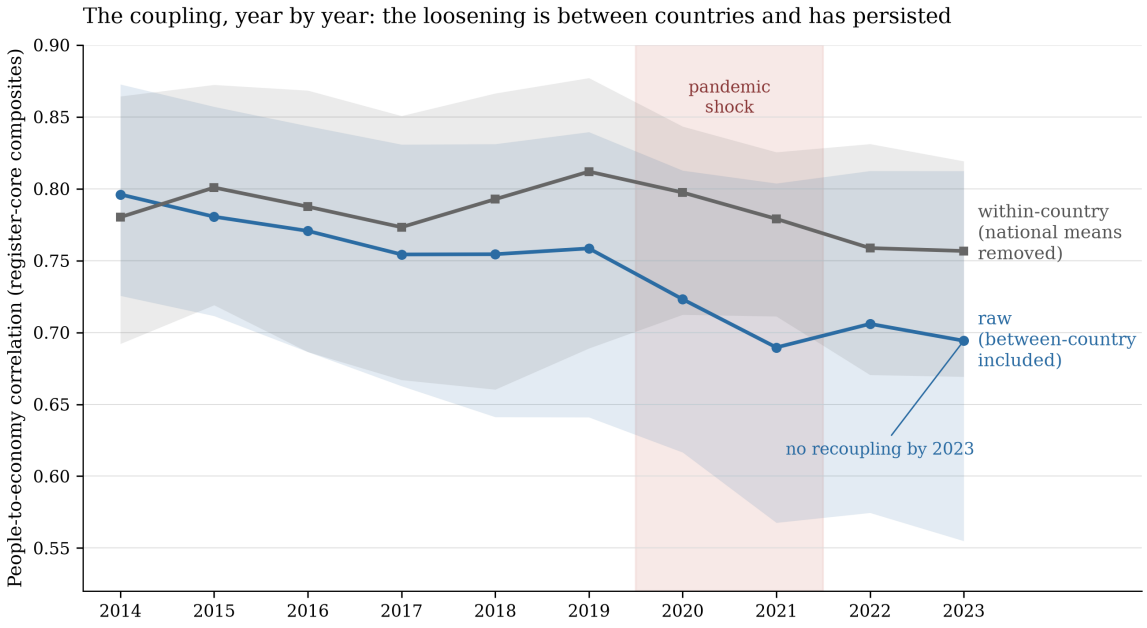

Fixed balanced panel of 139 NUTS 2 regions in 25 countries, register-only indicators, equal-weight composites re-standardised each year.
Bands: 95 per cent country-cluster bootstrap (1,000 draws).

The remaining three lines of evidence, individually weaker but consistent in direction, complete the picture.

- A second form of selectivity appears across tiers, and should not be conflated with the first.

The headline selectivity is across pillar pairs, the people-to-economy relation loosening while the place-based relations hold. Separately, a causal-forest estimate of the place-to-economy effect weakens most in the struggling and high-performing tiers across the shock, a pattern across tiers rather than across pairs, whose cross-period changes are not individually significant, it therefore enters as descriptive support rather than as a second test.

- The typology drifts downward across the shock, fifty regions descending and only twelve rising, predominantly from the high-performing tier into the stable core. The direction holds with the United Kingdom excluded, though its magnitude is inflated by the loss of UK data.
- The syndrome network retains its overall shape, with edge strengths correlated at 0.91 across the two periods, while its communities reshuffle, the place-based pillar is the most internally stable of the three under the shock.

Four methods thus point the same way: under the pandemic the people-side and the economy drew apart between national systems, while the place-based correlations and the within-country coupling held. Each signal is individually moderate, together they carry the framework's claim that regional development is a coupled system.

Section 5: Development paths, basins and traps

A typology shows where a region stands, but the deeper claim of the framework concerns how regions arrived there and why some cannot leave. Three snapshots cannot model these paths, although they can describe them. It is useful to picture a development landscape made up of valleys of differing depth, with a deep and persistent trap at one end and a shallow basin at the other from which regions readily slip, and with the pandemic tilting the whole surface so that regions descend far more easily than they climb.

A development landscape: deep basin at the bottom, shallow basin at the top

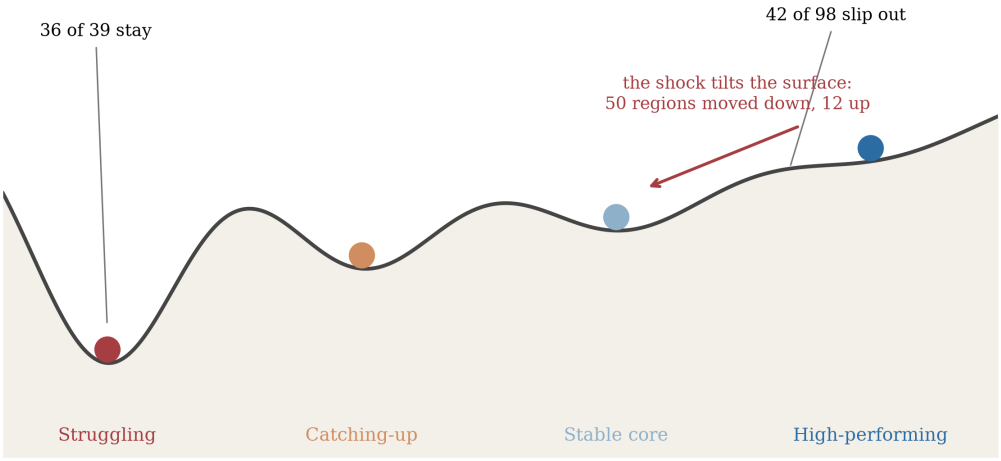

Schematic: the surface is illustrative, the annotations are the measured retention and transition counts (Appendix A.2, A.5).

- The tiers behave as basins of differing depth. Across the shock the struggling tier retains 92 per cent of its regions and the high-performing tier only 57 per cent, with the two intermediate tiers between, so the bottom of the ladder holds firm while regions leave the top with comparative ease.

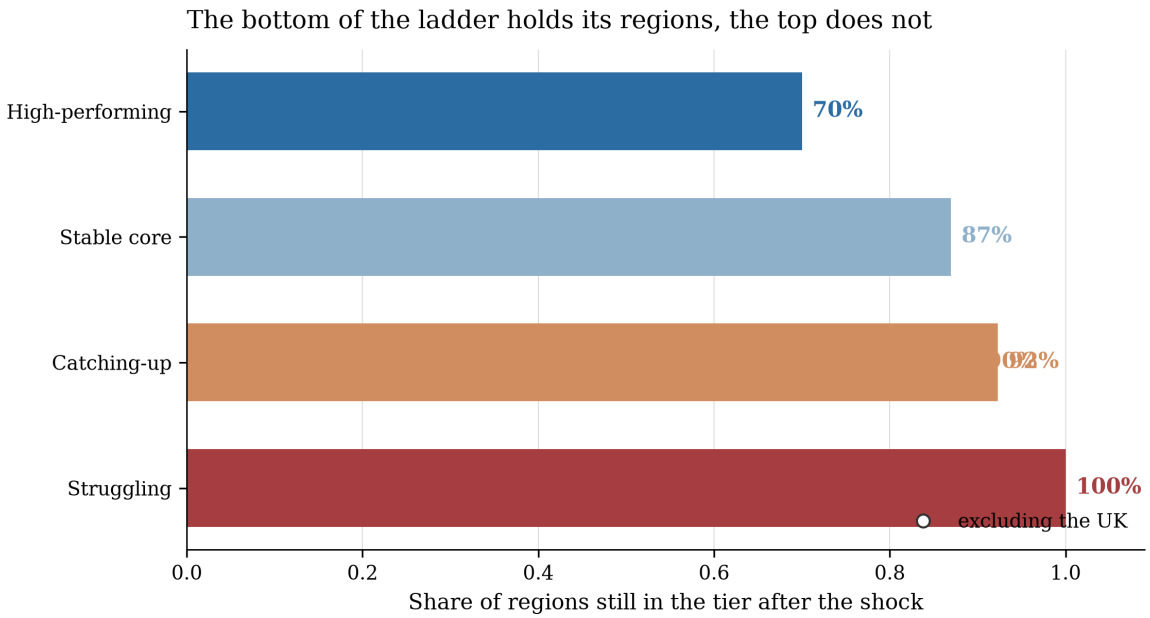

- The movement is markedly asymmetric. Downward transitions exceed upward ones by more than four to one, and the high-performing regions that slipped had been growing more slowly than those that remained. The obvious objection is mean reversion: a tier defined by high scores holds some regions that sit there partly on luck, and luck fades, so a leaky top might be mechanical rather than real. The calm-window transition answers this directly. Across the two pre-shock windows, 2014 to 2016 and 2017 to 2019, the high-performing tier retained 94 per cent of its regions and upward moves outnumbered downward ones by roughly two to one. Regression to the mean and measurement error are present in that window as well, so they can account for at most the six points of exit seen there, not the forty-three points seen under the shock. It is the shock that turns a gentle upward drift into a descent.

Where regions moved across the shock: 50 went down, 12 went up

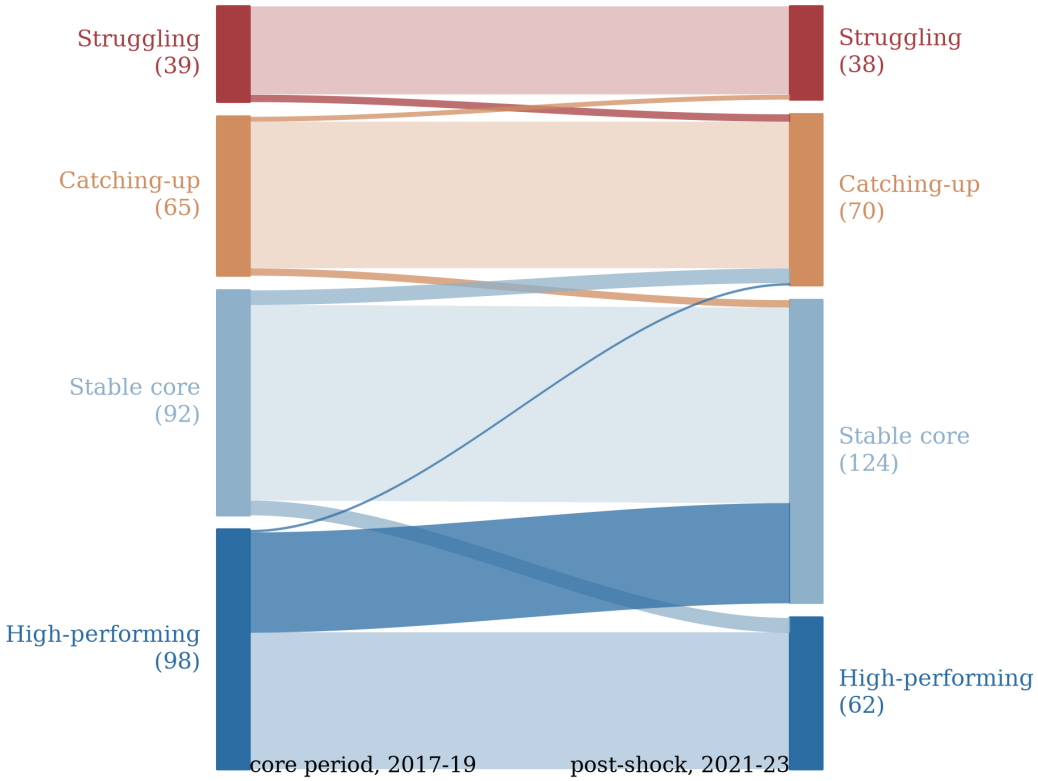

- This retention asymmetry is what makes the allostatic reading appropriate. An allostatic equilibrium is the most stable configuration a system can reach given the load it bears, and not necessarily a favourable one. A struggling region occupies a stable but impoverished equilibrium, which is why it holds its position so firmly and why expenditure that does not alter the underlying load does not move it. Allostasis here is an analogy anchored to the retention pattern, not a measured mechanism.
- The reading also connects to the typology literature. Static typologies (Iammarino, Rodriguez-Pose and Storper, 2019; Borsekova and Korony, 2023; Fratesi and Perucca, 2018) capture a region's position but say nothing about its trajectory. The four tiers align with the persistent development clubs of Iammarino, Rodriguez-Pose and Storper (2019), and the struggling tier corresponds to the development trap of Diemer et al. (2022).
- Two external comparisons come back mixed. The tiers track recomputed income clubs closely (Spearman 0.74 to 0.79), but because the tiers and the clubs share income data, this confirms only that the typology falls where the income literature would place it. A more independent test, reproducing the Diemer development trap, fails on a two-period growth proxy, since two pandemic-straddling periods cannot substitute for a comparison against a region's longer history. The trap therefore enters as lineage and as motivation for the modelling paper, not as a result reproduced here.
- The asymmetry carries a direct implication for policy. The two ends of the ladder are different problems. The struggling tier is a deep trap, so that funding which leaves the binding constraint untouched leaves the region where it is. The high-performing tier loses nearly half its members, so that the task there is to maintain the coupling which sustains it rather than to add further resource. The intermediate tiers move least of all. Effort distributed evenly along the ladder is therefore spent mostly where least is changing, and the diagnostic value of the typology is to

match the instrument to the tier, altering the load at the bottom of the ladder and preserving the coupling at the top. The stakes are heaviest at the bottom, where the trapped regions are also the places the discontent literature associates with long-run decline (Rodriguez-Pose, 2018; Diemer et al., 2022).

Read in this way, the typology is a set of basins of differing depth, anchored in the retention asymmetry and in the slower growth of the regions that descended. Its limit is equally clear: three snapshots can open these paths but cannot trace them along their length, which is the task of the modelling paper.

5.2 The SAGE decoupling in regional form

The people-to-economy loosening is the regional, post-pandemic counterpart of a decoupling already argued at the national and secular level.

- It extends the SAGE decoupling of Lima de Miranda and Snower (2020), with my earlier work (Conway, 2020) as the theoretical precursor that shows the dimensions of wellbeing separating from consumption under a shock. Here the loosening is shock-induced and measured from a high baseline, and it is a weakening of the link rather than its severance. The annual trace in Section 4 supplies the recoupling window that the two-period design lacked: no recoupling had occurred by 2023.
- The finding sits alongside evidence pointing the other way. Burger et al. (2025) find that subjective wellbeing raises regional productivity, a one-point higher wellbeing score corresponding to roughly 19 per cent higher total factor productivity across EU regions, and Veneri and Murtin (2019) document the divergence of income from broader living standards across OECD regions. Both are consistent with the framework rather than opposed to it: the pillars are coupled in normal times, which is what Burger and colleagues measure and what the pre-pandemic correlation of 0.85 records here, and the contribution of this paper is that the coupling loosens selectively under the shock. The coupling result thereby bounds the decoupling claim.
- The cross-pillar relations measured here have established precedents: Rodriguez-Pose and Ketterer (2020) on institutions and economic structure, Muringani et al. (2021) on social and human capital, Cortinovis et al. (2017) on formal and informal institutions, and Charron, Dijkstra and Lapuente (2014) on the path from institutions to trust.
- In terms of political economy, the two structural findings of this paper underlie a phenomenon the field already studies. A region, or a national system, held in a deep basin is the structural form of the places that do not matter, where decline is persistent rather than transient (Rodriguez-Pose, 2018), a shock that separates the people-side from the economy is, at the European scale, a failure of social and economic prosperity to recouple (Lima de Miranda and Snower, 2020), and these are the conditions under which the literature on discontent and anti-establishment voting situates the political reaction (Dijkstra, Poelman and Rodriguez-Pose, 2020; Cavallaro and Zanetti, 2020). Appendix A.10 checks this connection directly against the typology. In the last national election before the pandemic, anti-establishment parties took roughly 40 per cent of the vote in the struggling and catching-up tiers, against 16 to 20 per cent in the stable core and high-performing tiers, and the gradient survives with country means removed, so it is not a country composition effect. The composition of that vote also differs along the ladder, with the far right strongest in the catching-up tier while the discontent of the struggling tier flows more to the populist left. This is a descriptive cross-section, not a test of the dynamic from trap to discontent, the direct demonstration that trap and discontent coincide remains Rodriguez-Pose, Dijkstra and Poelman (2024). The contribution here is the

coupled-ecosystem structure beneath those same regions, the deep basins and the people-to-economy decoupling that the discontent literature takes as given, and the integration of that structure with their mapping is the task the modelling paper undertakes.

The decoupling is therefore a bounded, shock-induced, regional instance of an established macroeconomic pattern, and it carries real political stakes through the connection from development traps (Diemer et al., 2022) to the geography of discontent (Dijkstra, Poelman and Rodriguez-Pose, 2020).

Section 6: The pairwise literature within one structure

The conceptual companion began from a complaint: the regional-development literature establishes the parts of the system two at a time and never assembles them. Within this panel, the assembled version can be checked directly.

- The procedure is simple: for every pair of indicators that crosses a pillar boundary, the raw correlation across regions, which is what a pairwise study reports, is compared with the correlation that remains once a single general development factor has been removed. The first asks whether the literature's relations appear in the data, the second asks whether each is a direct connection or merely a reflection of the shared development dimension.
- The relations do appear, and in the expected direction. Government quality and trust, social capital and skills, institutions and the economy, infrastructure and innovation all emerge as strong raw correlations, between 0.4 and 0.8, so the relations the literature reports are all present in this panel.
- Most of them, however, reflect the shared development dimension. Across the 224 strong cross-pillar correlations, removing the single development factor reduces the average from 0.50 to 0.18 and sends more than a third of them to near zero. Much of what the field has established, one pair at a time, is the same underlying development dimension appearing in different guises, which is the one-factor result of Section 2.2 seen from the side of the literature.
- A specific web of relations nonetheless survives beneath this shared dimension. Fifty-six relations between distinct domains remain as direct connections once everything else has been conditioned out, forty-seven of them even after the general factor is removed, and roughly nine in ten persist across the shock. Among them are household income and adult learning, clean air and social contact, new business formation and job-skill matching, and government quality and civic participation.

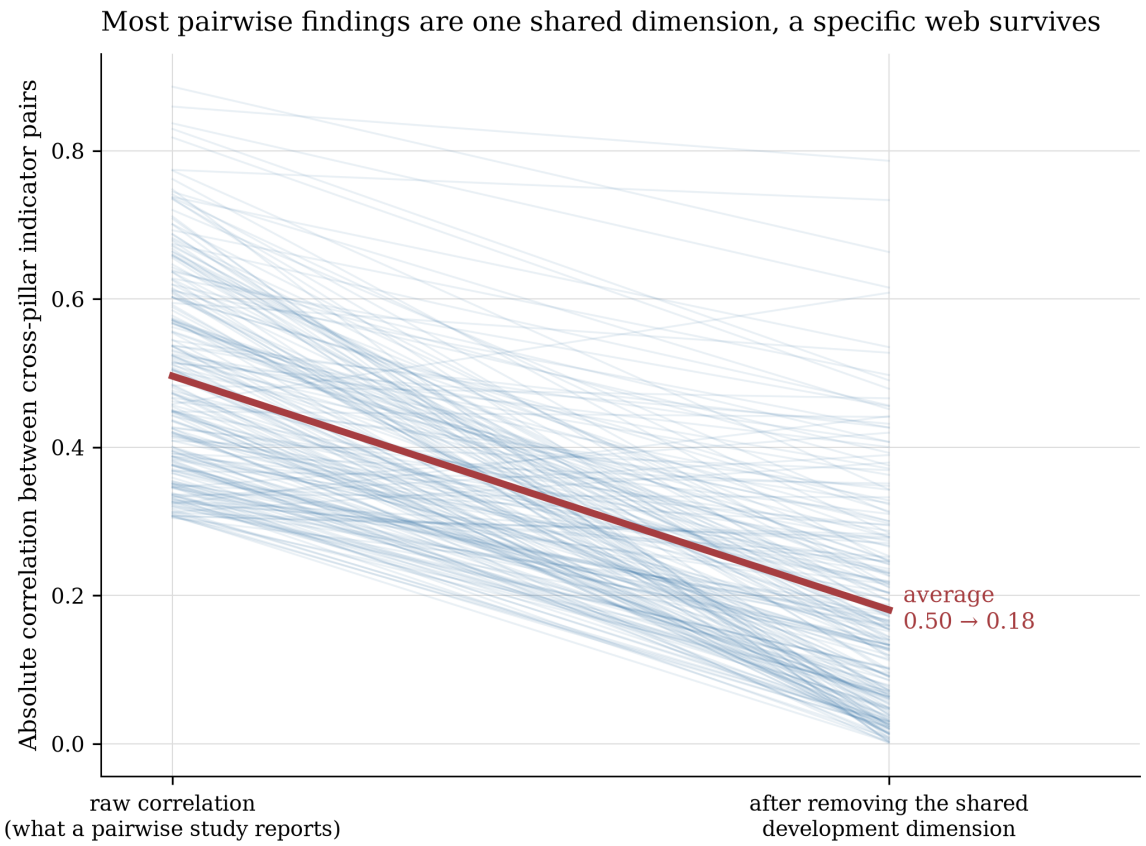

- The relations that survive are interpretable. The two findings from the literature that hold as direct connections, rather than as reflections of the shared dimension, both concern institutions reaching into the social fabric, government quality to civic participation and to institutional trust, the Charron and Cortinovis pathways, whereas government quality to the economy and social capital to growth prove to be largely a reflection of it.
- A caution is in order. These are descriptive partial correlations, implying neither direction nor causation, the claim rests on the overall pattern rather than on any individual small relation, and the surviving web is cross-pillar, so that it no more reconstitutes the three pillars than the general factor does. This is consistent with the reading of the pillars as a conceptual organisation given in Section 2.2.

On this account each study in the literature was correct in what it found, but none was complete, since each observed a single relation and none observed the whole. The ecosystem measurement recovers these relations, separates the shared development dimension from the specific connections, and presents them as one structure. This is the integrative contribution of the approach. The evidence that the system is real is the shock of Section 4, what this section adds is that the structure the shock reveals also contains the field's established results, which returns the argument to the opening problem of the conceptual companion.

Section 7: Conclusions

Rather than resting on a single decisive result, the paper assembles a multi-method account of the European regional economy as a coupled system.

- Regional development is strongly and spatially organised (Section 2.1), the three pillars organise it as a conceptual structure and a reliable composite, though they are not three separable dimensions (Section 2.2), the framework yields a robust four-tier ladder and a layer of syndromes (Section 3), and the structure responds to the pandemic selectively, in a manner that aggregate measures do not detect (Section 4).
- This structure is organised predominantly at the national level: between three-quarters and nine-tenths of the development gradient is between-country, the typology is a cross-national ladder observed at regional resolution, and the decoupling under the shock is a divergence between national systems while within-country coupling is maintained (Section 3 and Appendices A.8 and A.9).
- Strong coupling in calm times combined with selective decoupling under stress is what the framework means by interconnection, and the combination separates a coupled system from both a single factor and a collection of independent parts.
- Understood as basins of differing depth (Section 5), the typology forms the bridge to the dynamic model, in which the tiers become multiple equilibria and the question becomes what retains a region within a basin and what displaces it.
- Finally, the fragmented pairwise literature is recovered as one structure (Section 6), most of it the reflection of a single development dimension with a stable web of specific relations beneath, which returns the empirical paper to the opening problem of the conceptual companion.

The framework emerges from the test as what the conceptual companion says it is: a coupled system that shows its structure under stress, settles into configurations that are stable but can shift, and runs primarily through national systems with regional texture beneath them. Two results become foundations for the modelling paper that follows: development is organised through national systems that carry within-country regional structure, and the regions within them sit in basins of differing depth. The model takes these as its starting conditions, a population of heterogeneous regions nested within national systems, and asks what holds a region in its basin and what tips it out, the question this paper can pose but not answer.

References

Borsekova, K. and Korony, S. (2023). Resilience and vulnerability of regional labour markets: principal component analysis of labour market efficiency in the EU. Regional Studies, 57(12), 2373 to 2390.

Burger, M. J., Courchesne, S., Kostarakos, I., Peroni, C., Sarracino, F. and Stavropoulos, S. (2025). Subjective well-being and regional productivity in the European Union. Annals of Regional Science, 74(3).

Cavallaro, M. and Zanetti, M. A. (2020). Divided we stand: attitudes, social classes, and voting for the radical right after the Great Recession in Europe. Ethnic and Racial Studies, 43(2), 313 to 332.

Charron, N., Dijkstra, L. and Lapuente, V. (2014). Regional governance matters: quality of government within European Union member states. Regional Studies, 48(1), 68 to 90.

Conway, A. (2020). Wellbeing and Macroeconomics: a SAGE approach. Master's thesis, Sciences Po Paris.

Conway, A. (in preparation). A Regional Economic Development Ecosystem Framework: Integrating Place, People, and Economy in European Regions.

Cortinovis, N., Xiao, J., Boschma, R. and Van Oort, F. G. (2017). Quality of government and social capital as drivers of regional diversification in Europe. Journal of Economic Geography, 17(6), 1179 to 1208.

Diemer, A., Iammarino, S., Rodriguez-Pose, A. and Storper, M. (2022). The regional development trap in Europe. Economic Geography, 98(5), 487 to 509.

Dijkstra, L., Poelman, H. and Rodriguez-Pose, A. (2020). The geography of EU discontent. Regional Studies, 54(6), 737 to 753.

Epskamp, S. and Fried, E. I. (2018). A tutorial on regularized partial correlation networks. Psychological Methods, 23(4), 617 to 634.

Fratesi, U. and Perucca, G. (2018). Territorial capital and the resilience of European regions. Annals of Regional Science, 60(2), 241 to 264.

Haslbeck, J. M. B. and Waldorp, L. J. (2020). mgm: estimating time-varying mixed graphical models in high-dimensional data. Journal of Statistical Software, 93(8), 1 to 46.

Iammarino, S., Rodriguez-Pose, A. and Storper, M. (2019). Regional inequality in Europe: evidence, theory and policy implications. Journal of Economic Geography, 19(2), 273 to 298.

Lima de Miranda, K. and Snower, D. J. (2020). Recoupling economic and social prosperity. Global Perspectives, 1(1), 11867.

Muringani, J., Fitjar, R. D. and Rodriguez-Pose, A. (2021). Social capital and economic growth in the regions of Europe. Environment and Planning A, 53(6), 1412 to 1434.

Rodriguez-Pose, A. (2018). The revenge of the places that don't matter (and what to do about it). Cambridge Journal of Regions, Economy and Society, 11(1), 189 to 209.

Rodriguez-Pose, A., Dijkstra, L. and Poelman, H. (2024). The geography of EU discontent and the regional development trap. Economic Geography, 100(3), 213 to 245.

Rodriguez-Pose, A. and Ketterer, T. (2020). Institutional change and the development of lagging regions in Europe. Regional Studies, 54(7), 974 to 986.

Rooduijn, M., Pirro, A. L. P., Halikiopoulou, D., Froio, C., van Kessel, S., de Lange, S. L., Mudde, C. and Taggart, P. (2024). The PopuList: a database of populist, far-left, and far-right parties using expert-informed qualitative comparative classification (EiQCC). British Journal of Political Science, 54(3), 969 to 978. Version 4.0 of the database used here.

Schraff, D., Vergioglou, I. and Demirci, B. B. (2023). The European NUTS-level election dataset: a tool to map European electoral geography. Party Politics, 29(3), 570 to 579.

Veneri, P. and Murtin, F. (2019). Where are the highest living standards? Measuring well-being and inclusiveness in OECD regions. Regional Studies, 53(5), 657 to 666.

Appendix A: Test ledger

Every test run, what was asked, what was returned, and the verdict carried into the paper. Grading is conservative, 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 across all four specifications, 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 0.195 to 0.217, loose by design to preserve indicator diversity	Documented limitation
One-factor shadow (Phase 1.3)	Is the typology just the general factor?	One-factor four-cluster vs published typology ARI 0.89 to 0.93, 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	Are the tiers reproducible?	Core 0.83 to 0.91, post-shock 0.86 to 0.99, all above 0.75	Robust
Centroid ordering	Ladder or configurations?	Centroids strictly ordered on all three pillars (-3.98, -1.57, +0.49, +2.17)	Ordered ladder, configurational claim dropped
Cross-period persistence	Is the typology stable across the shock?	78.9 per cent stay (232 of 294), ARI +0.515	Stable

Test	Question	Outcome	Verdict
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 to 0.705 vs random near 0.28	Contiguous

A.3 Coupling and the shock (Section 4 signals)

Test	Question	Outcome	Verdict
Composite decoupling (Stage 4)	Does People to Economy loosen?	0.847 to 0.776, Place to Economy only 0.032, Place to People only 0.029, Wald z = 2.51, p = 0.012	Selective, signal one
Country block bootstrap (Phase 1.1)	Does it survive clustered inference?	Direction and magnitude robust, design effect 0.99, jackknife all positive, borderline at strict joint standard (Bonferroni p 0.051 to 0.058)	Robust in direction, borderline in significance
Complete-case (Phase 1.2)	Is it an imputation artefact?	Persists 0.793 to 0.586 on 124 fixed complete cases	Not an artefact, coverage caveat stated
Furlough-robust subset	Is it a labour-flow measurement artefact?	Survives dropping employment, activity, unemployment and long-term unemployment (0.788 to 0.739), sharper on the output, productivity and income core (0.720 to 0.623), observed data, no imputation	Not a furlough artefact, strongest in the least distortable indicators

Test	Question	Outcome	Verdict
Directional and pre-trend reading	Does significance rest on a borderline p-value?	People to Economy flat across pre-pandemic and core windows (0.79 to 0.83, 0.78 to 0.82), falls to mid-0.6s only post-shock, directional test of a stated prior	Carried by the before-after pattern, not the threshold
Annual coupling trace (register core)	Does the window contrast hold year by year, and does the link recouple?	Balanced 139-region panel, 25 countries, 2014 to 2023: raw People to Economy drifts 0.80 to 0.76 over 2014 to 2019, steps to 0.69 by 2021, still 0.69 in 2023, within-country holds 0.76 to 0.81 throughout	Trajectory confirms the contrast, no recoupling by 2023
Causal forest CATE (Stage 6)	Does Place to Economy weaken by tier?	Concentrated in struggling and high-performing tiers, cross-period changes not individually significant (p 0.41 to 0.78)	Descriptive pattern, signal two
Typology drift (Stage 3)	Does the ladder slip downward?	50 down vs 12 up, survives UK exclusion, magnitude inflated by UK attenuation	Directional, signal three
Network temporal stability (Stage 2)	Stable fabric or reshuffle?	Edge-weight r 0.91 on shared edges, community ARI 0.19, Place most shock-stable	Stable fabric, reshuffled syndromes, signal four
Pre-pandemic coupling baseline	Is high coupling the calm-time norm?	People to Economy about 0.79 to 0.83 across pre-pandemic windows	Confirms calm-time coupling

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 core, six post-shock, ARI vs pillar partition at most 0.121	Cross-cutting syndromes
Syndrome strength	Stronger than chance?	5.5 to 8.3 times stronger than 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
Diemer trap (Phase 1.5)	Does struggling equal the development trap?	Rank correlation near zero on a two-period growth proxy	Null, data limitation, not refutation
Basin retention (Phase 1.5)	Are tiers basins of differing depth?	Retention struggling 0.92, top 0.57, leaky-top strengthens excluding UK	Clean support for basins
Calm-window placebo (2014-16 to 2017-19)	Is the leaky top just regression to the mean?	Non-shock transition: persistence 86 per cent, high-performing retention 0.94, upward moves outnumber downward about 2 to 1, against shock persistence 79 per cent, top retention 0.57, downward about 4 to 1	Mechanical exit at most 6 points, the leaky top is shock-specific
Global Moran's I (Stage 5)	Are pillars spatially clustered?	Core 0.84 to 0.90, post-shock 0.80 to 0.86, 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 decoupling is moderate per signal (5 to 6 out of 10) and stronger only as a convergent four-signal story (7 to 8 out of 10). No single test is decisive.
- The three-pillar structure is not statistically separable at the indicator level, the framework's value is reliability and downstream outputs, not latent distinctness.
- Two periods cannot identify dynamics. Basins, traps, and the allostatic regulation are described and motivated, not modelled. Left to Paper 3.
- The development-trap reproduction is null on the available proxy, the concept is carried as lineage, not as a reproduced result.
- UK post-Brexit attrition inflates several magnitudes, every UK-sensitive result is reported with its UK-excluded counterpart.
- National context dominates the level of development (Appendices A.8 and A.9): between three-quarters and four-fifths of the gradient is between-country, the typology is a cross-national ladder with within-country texture, and the shock-period decoupling is a between-country divergence while within-country coupling holds. The structure and the decoupling are therefore national-regime phenomena read at regional resolution.
- The framework is structural and describes configurations, the political economy that produces and holds them, partisan governments, welfare regimes and redistributive choices, sits outside the object as a deliberate scope choice, distinct from the dynamic regulation left to Paper 3.
- 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 two-period regional snapshot cannot capture and sit outside the design.
- The link from trapped regions to political discontent is checked descriptively in 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. The check is cross-sectional and uncontrolled, the dynamic from trap to discontent stays with the literature (Rodriguez-Pose, 2018; Diemer et al., 2022) and with the modelling paper.

A.7 Integration with the pairwise literature (Section 6)

Test	Question	Outcome	Verdict
Marginal reproduction	Do the literature's cross-pillar couplings appear in our data?	All reproduce as raw correlations of 0.4 to 0.8 in the expected direction	Connected to the named findings

Test	Question	Outcome	Verdict
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
Direct backbone	Is there structure beyond the tide?	66 direct cross-pillar edges, 56 genuine cross-domain after removing 10 same-construct duplicates, 47 survive removing the general factor	A stable specific web remains
Cross-period persistence	Is the web stable?	About 89 per cent of direct cross-pillar edges persist from core to post-shock	Stable, not noise
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 (84 per cent post-shock), within-country 12 to 16 per cent	National context sets the level

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 (60 per cent post-shock), several span three, average two tiers per country	The typology resolves within-country heterogeneity
Locus of the decoupling	Within-country or between-country?	Within countries the people-to-economy correlation is flat to rising across the shock (0.65 to 0.76), the raw fall is entirely between-country	The decoupling is a between-country divergence, not a within-country one

A.9 Robustness to the indicator set

Three pre-specified operationalisations, reported 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. Composite-level recomputation on observed data, a full pipeline re-run per set is available but not required for these headline checks.

Operationalisation	Between-country share	People-Economy, raw (core to post-shock)	People-Economy, within-country	Typology agreement (ARI vs headline)
Full set	83%	0.81 to 0.69	0.65 to 0.71	0.57
No survey or government-quality	79%	0.85 to 0.71	0.73 to 0.76	0.50
Register-only core	74%	0.85 to 0.71	0.73 to 0.76	0.41

The headline findings hold across all three: national context dominates the level (74 to 83 per cent between-country), the raw people-to-economy decoupling appears in every set, the within-country coupling holds or tightens in every set, 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, not instability.

A.10 The typology and the geography of discontent

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. 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
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, the regional layer carries part of it.
- The far right alone is not monotone. 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.
- Caveats. Vote shares are measured at one election per country, the party classification is national, the table is descriptive, with no controls. It corroborates the association of the lower tiers with anti-establishment voting (Dijkstra, Poelman and Rodriguez-Pose, 2020; Rodriguez-Pose, Dijkstra and Poelman, 2024) at the resolution of this paper's typology. It does not test the dynamic from trap to discontent.

Appendix B: Figure and data provenance

- All 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`), syndrome network (`networks/core_pcor_pooled.parquet`, `networks/core_communities.parquet`), and basin and transition statistics (`clusters/basin_descriptives` `clusters/cross_period_transitions.parquet`).
- Maps use Eurostat NUTS boundary files (1:1 million), 2016 vintage for the core period and 2021 for the post-shock period, 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 consistently across both panels. Map join coverage: 287 of 294 regions core, 289 of 294 post-shock, unmatched regions are vintage boundary changes, not 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, the post-shock panel uses Hungarian-aligned labels so the colours mean the same thing in both periods.
- Figures: `fig_typology_maps.png` (Section 3), `fig_network.png` (the partial-correlation connectogram, Section 2.2), `fig_decoupling_map.png` (the people-minus-economy residual map, Section 2.2), `fig_decoupling.png` and `fig_decoupling_scatter.png` (Section 4), `fig_basins_landscape.png` (the development-landscape schematic, Section 5), `fig_basins.png` (Section 5), `fig_transitions.png` (the alluvial flow, Section 5), and `fig_tide_wiring.png` (the tide-versus-wiring slope chart, Section 6). Every figure is exported twice from one script (`code/analysis/fig_paper2_figures.py`, with the annual trace in

`fig_coupling_trace.py`): high-resolution PNG for the working draft and 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 development-landscape figure is a schematic, not a data plot: the undulating surface is illustrative, but the annotations are exact, the four wells are the four tiers ordered by the general development factor, with the retention and transition counts from Appendix A.2 and A.5 (36 of 39 struggling regions stay, 42 of 98 high-performing regions slip out, 50 regions move down against 12 up). The deep narrow well and the shallow open basin encode the retention asymmetry, the lowered ridge and the arrows encode the shock-driven downward drift.
- The annual coupling trace (`fig_coupling_trace.png`, Section 4) is built from the raw Eurostat annual series in the analysis repository (`data/raw/eurostat/`): register-only indicators (tertiary attainment, early leavers, adult learning, NEET, GDP per head in PPS, labour productivity, employment rate, unemployment rate), equal-weight sign-adjusted composites re-standardised each year on a fixed balanced panel of 139 NUTS 2 regions in 25 countries, 2014 to 2023. Median income (ilc_di04) was excluded for lack of NUTS 2 coverage. Raw and country-demeaned Pearson correlations per year, 95 per cent bands from a country-cluster bootstrap (1,000 draws), series stored in `sem/coupling_trace_annual.parquet`.
- 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 transition figure is an alluvial flow in which ribbon width is the number of regions and ribbon colour is the origin tier, so the downward dominance is read directly.