

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

Conway

Contents

Abstract	1
Section 1: Introduction	2
Section 2: Literature	3
Section 3: Data and methods	5
Section 4: Results	8
Section 5: Discussion	23
Section 6: Conclusions	26
References	27

Abstract

In this paper I test the ecosystem framework of the conceptual companion, which organises regional development around three interacting pillars: Place-based Conditions, Human and Social Capital and Economic Activity. I use a cross-section of thirty-five indicators for 294 NUTS 2 regions in 2017 to 2019 and an annual panel of 152 regions in 23 countries from 2008 to 2023. Development is strongly organised in space, a single general factor fits the indicators better than three, and regions sort into four stable tiers, with between two thirds and nine tenths of the gradient lying between countries. The two shocks in the panel behave differently. Between 2008 and 2013 Mediterranean regions fell 0.52 standard deviations in relative position, with the economic pillar taking almost all of the fall and still displaced a decade later, whereas in the pandemic years no bloc moved by more than 0.18 and the South gained. The regions worst hit in one episode were no more likely than chance to be worst hit in the other. The correlation between the people and economy composites has declined since 2011, a slide that began with the recovery from the euro crisis. A development-trap index built to the standard specification agrees with the two-shock account by bloc but bears no relation to the tiers. Within countries anti-establishment voting follows a region's level, chiefly its income, at the last election and more weakly at the one before, and bears no relation to its years in the trap, with both measured before each vote. A partial-correlation network over the indicators shows that, in this panel, most of the association behind the pairwise relations reported in the literature lies on one shared development dimension. The analysis is descriptive, and mechanisms are left to the modelling companion.

Section 1: Introduction

1.1 What the paper asks

The conceptual companion sets out the ecosystem framework as a proposal, in which regional development is a coupled system instead of a stack of separate parts. It does not test that proposal, and this paper does. I ask whether the framework holds in the European data, and what a region looks like as a coupled system when it is put under stress.

The lens is wellbeing in the regional economy. I therefore look past output to the people-side of a region and the territorial conditions beneath it, 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 a window onto the social side of development and not a direct measure of how people feel, so any loosening documented here is a loosening between measured conditions.

The unit is the European region at NUTS 2 level, which is the scale at which Cohesion Policy operates and the scale at which all three pillars can be observed together.

I treat the framework as a set of testable commitments and not as a metaphor. The commitments are three organising domains, named mechanisms running between them, and the claim that the system is coupled and settles into configurations that are stable and able to shift.

A coupled system shows itself under stress, so the design needs a shock. I first built the design around the pandemic, which proved the less informative of the two shocks available. Extending the panel back to 2008 brings in the euro crisis, which was larger, arrived through different channels and left a mark still visible fifteen years later. With two shocks in the panel I can say which features of a regional system are properties of the region and which are properties of the shock.

1.2 The organising idea

The paper has three main results, and they need to be read together.

In calm conditions 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 euro crisis they came apart. In southern Europe the economic pillar fell 0.82 standard deviations relative to the rest of the continent between 2008 and 2013, while the human and social pillar fell 0.19 and the place-based pillar rose 0.10. Eleven years later the economy was still displaced, whereas the other two had returned.

Under the pandemic nothing came apart. The regions hurt by the first shock were not those hurt by the second, since of 152 regions none was in the worst quarter of both.

Three separable dimensions would not produce the first result, a single underlying dimension with nothing specific to each pillar could not produce the second, and an account in which some regions are fragile whatever the shock cannot produce the third. The account that fits all three is a coupled system whose parts separate when a shock arrives through a channel that reaches one of them first, and whose vulnerability at any moment depends on where the shock comes from. A common factor with pillar-specific shocks would fit them as well. It differs from a coupled system only in whether

an imbalance between the pillars is later corrected, which the modelling companion estimates. On its panel a model with the links between pillars fits better, by more than eighty log-posterior units, than one with common national shocks and no links.

For this reason I read the single-factor result as evidence for the framework, although at first sight it looks like a refutation, because a coupled system should produce high correlation in calm times.

1.3 Contribution

The paper makes three contributions, all drawn from the same regions and the same three composites.

First, I test the three domains of regional development together, across two shocks, on an annual panel. The pairwise literature of Section 2.3 has connected the domains two at a time, and the measurement traditions have scored them side by side. Testing them together across sixteen years locates the change and separates the properties of regions from the properties of shocks. The euro crisis reordered Europe's regions along national lines and took more than a decade to work off, whereas the pandemic did not reorder them at all.

Second, I follow the same regions year by year over sixteen years, which the typology literature, working from snapshots, has not done, and I bring the field's own instrument to bear on its own question. A development-trap index built to the specification of Diemer, Iammarino, Rodríguez-Pose and Storper (2022) agrees with the two-shock account about which part of Europe lost a decade, yet bears no relation to the four tiers. A low position on the ladder and a stalled trajectory are therefore separate conditions, and only one of them tracks anti-establishment voting within countries.

Third, I use a method that is rare in regional studies, a partial-correlation network estimated over the development indicators themselves, to split what the pairwise literature has found into one shared development dimension and a smaller set of direct links beneath it.

The paper neither proposes a new composite index nor ranks regions for policy purposes. The composites are instruments for testing the framework, and their high correlation is one of the results. The analysis is descriptive, in the tradition of the clubs and traps work it builds on, and it measures the structure of European regional development and how it moved under two shocks. Identifying what causes what, and measuring how people feel, lie beyond its design.

The evidence has two limits. First, much of the variation between regions is the European east-west income gap, so the national finding describes economic geography as much as institutions. Second, the annual panel excludes the United Kingdom, Norway, Switzerland and Turkey. It measures the people-side through education and health, since trust and participation are available only in survey rounds.

The remainder of the paper proceeds as follows. Section 2 reviews the literature the results speak to. Section 3 sets out the data and methods, including how the typology is built and how the annual panel differs from the wide panel. Section 4 reports the results, Section 5 discusses them and what they imply for Cohesion Policy, and Section 6 concludes.

Section 2: Literature

The results of this paper speak to four strands of work: regional typologies and development clubs, together with the evidence on how European regions fared in the crisis before the pandemic, the de-

coupling of wellbeing from wealth, the evidence on interactions between the domains of development, and the geography of discontent.

2.1 Regional typologies, clubs, traps and the crisis

A long line of work classifies European regions into types. Iammarino, Rodríguez-Pose and Storper (2019) document the formation of development clubs with divergent and self-reinforcing trajectories, and Borsekova and Korony (2023) examine the resilience and vulnerability of regional labour markets through a principal component analysis of labour market efficiency. Borsekova et al. (2026) sort 233 regions by decision trees over institutional and innovation indicators, and find strong institutions to be a precondition for innovation-led growth. Fratesi and Perucca (2018) relate territorial capital to the resilience of European regions. Diemer et al. (2022) identify a regional development trap, established on decades of data, in which a region is unable to retain its economic dynamism relative to its own past and to its national and European peers. Static typologies of this kind capture a region's position well, but they say little about how regions move between types, and this is where the paper adds to them.

A smaller literature examines how European regions fared in the crisis that preceded the pandemic, and it frames the question differently, in a way that matters here. Dijkstra, Garcilazo and McCann (2015) find the contraction concentrated in urban regions and in remote rural ones, with intermediate and near-city rural regions more resilient, and note that in several countries the crisis inverted the pre-crisis advantage of the capital metro region. They read incidence by settlement type, whereas Crescenzi, Luca and Milio (2016) read it by nation, relating the depth of the regional contraction to pre-2008 national macroeconomic conditions across EU27 regions and finding a limited set of macroeconomic factors doing most of the work. Both study the depth of the initial contraction. Neither follows the displacement to see how long it lasted, and neither has a second shock to compare it with, which are the two things Section 4.4 adds.

These crisis studies belong to the wider literature on regional economic resilience, which asks how regions absorb a recessionary shock and whether they return to their previous path (Martin, 2012; Fingleton, Garretsen and Martin, 2012; Fratesi and Rodríguez-Pose, 2016; Crescenzi, Luca and Milio, 2016). The two-shock design of this paper is a resilience design in that sense. On this evidence, the overlap result of Section 4.4 indicates that these measures give no stable ranking of regions by resilience across two shocks that arrived through different channels.

2.2 The decoupling of wellbeing from wealth

Lima de Miranda and Snower (2020) show that across developed countries over 2007 to 2017, measures of social affiliation and individual empowerment moved apart from GDP per capita, a decoupling of social from economic prosperity that motivates the conceptual companion. In my earlier work (Conway, 2020) I found the same separation inside a heterogeneous-agent model, where the dimensions of wellbeing separate from consumption under a shock. Veneri and Murin (2019) document the divergence of income from broader living standards across OECD regions. Burger et al. (2025) point the other way, finding that subjective wellbeing raises regional productivity, with a one-point higher wellbeing score corresponding to roughly 19 per cent higher total factor productivity across EU regions. The two sets of findings are compatible, since the second measures the coupling of wellbeing and the economy in normal times and the first a separation that opens over time or under shocks.

However, the decoupling has so far been argued mostly at the national and secular level, and not for regions under a common shock.

2.3 Interactions between domains

A separate body of empirical work models the interactions between domains directly, one pair at a time. Rodríguez-Pose and Ketterer (2020) find that improvements in institutional quality drive growth in lagging regions. Muringani, Fitjar and Rodríguez-Pose (2021) find the economic return to bridging social capital conditional on a region’s human capital. Cortinovis et al. (2017) find that formal and informal institutions jointly drive regional diversification. Charron and Rothstein (2018) find the quality of regional government to be the strongest determinant of differences in generalised trust between regions of the same country. Each of these establishes a cross-domain relation while holding the rest of the regional system fixed. None of them can therefore tell how much of its relation is specific to the pair and how much reflects the general level of development that the regions share.

2.4 The geography of discontent

Rodríguez-Pose (2018) links long-run economic decline to anti-establishment voting in what he calls the places that don’t matter, and Dijkstra, Poelman and Rodríguez-Pose (2020) map the geography of discontent with the European Union across its regions. Cavallaro and Zanetti (2020) examine the attitudes and social classes behind voting for the radical right after the Great Recession. Rodríguez-Pose, Dijkstra and Poelman (2024) connect this geography to the development trap, showing that trap and discontent coincide across European regions. This literature takes the depth of a region’s decline as its starting point. The coupling between the social and economic sides of those regions, which this paper measures, sits beneath it.

Taken together, these strands leave three questions that the panels can address. The first is whether the domains that the pairwise studies connect two at a time behave as one coupled system when observed together, how that system responds to a shock, and whether it responds to two shocks in the same way. The second is how regions move between the types that the typology literature describes, and whether being low on the ladder is the same condition as being trapped. The third is how much of the pairwise evidence is specific and how much reflects the shared level of development.

Section 3: Data and methods

3.1 Data

The evidence rests on two harmonised panels of European regions that I assembled for this purpose. The first is wide enough to see all three pillars at once and is read as a single cross-section before the pandemic, while the second is narrower in what it measures and longer in time, reaching back across the euro crisis.

The two panels are different objects, and one figure in particular should not be equated between them. The correlation of 0.85 between Human and Social Capital and Economic Activity reported in Section 4.1 is computed between PCA composites of the thirty-five indicators across 294 regions in 2017 to 2019. The correlation of 0.85 reported for 2011 in Section 4.5 is computed between equal-weight pillars of the fifteen annual indicators across the 152 regions of the annual panel. They rest

on different regions, different indicators and different weights, and their agreement in level is not a replication of one by the other.

The wide panel covers 294 NUTS 2 regions and thirty-five indicators spread across the three pillars, in a core window before the pandemic (2017 to 2019), with an earlier window (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, drawn from the 2013, 2016 and 2021 NUTS vintages.

The indicator set is theory-led, in that each indicator, drawn from the standard regional sources, is there to measure one of the three pillars. The set is then trimmed for coverage, dropping anything too sparse to impute, and for redundancy. There is also a practical constraint, in that the multivariate methods need a single fixed, common and complete-enough matrix and cannot take a different set of indicators region by region. 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. Appendix A.9 shows that 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 multiple imputation by chained equations with random forests (van Buuren and Groothuis-Oudshoorn, 2011), with twenty imputed datasets pooled by Rubin’s rules (Rubin, 1987). Of the cells, 8.5 per cent in the core period are imputed. Indicators enter at the resolution their source provides. Across the 253 non-UK regions, 57 per cent of cells are direct NUTS 2 observations, 7 to 8 per cent are NUTS 1 values applied to their child regions, 24 to 26 per cent are national values applied to every region of the country, and 11 per cent are imputed. Ten of the thirty-five indicators are majority national values, including median income, inequality, deprivation, overcrowding, housing overburden and tertiary enrolment. National values carry no within-country variation by construction, which raises the between-country share reported in Section 4.3 and Appendix A.8, and removing those ten indicators lowers it from 0.87 to 0.81. However, they do not inflate the correlations between pillars, since removing them raises those slightly.

The channel variables of Section 4.4, employment by sector in 2008 and 2019 and tourist nights in 2019, come from Eurostat’s regional accounts and tourism statistics at NUTS 2 for all 152 regions, and are used only there.

The annual panel is built so that none of these resolution questions arises. It holds 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 NUTS 1 aggregate. Six measure Economic Activity: GDP per head, the employment, activity, unemployment and long-term unemployment rates, and employment in high-technology sectors. Six measure Human and Social Capital: tertiary attainment, the low-education share, early leavers from education, adult learning, young people not in employment, education or training, and infant mortality. Three measure Place-based Conditions: life expectancy at birth, road deaths and population density. 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. The balanced-panel rule drops regions whose codes changed during the window, which is why Poland contributes four of its seventeen regions. The United Kingdom is excluded because Eurostat ceased publishing its regional series, Turkey because it is an outlier on several indicators, and Norway and Switzerland because they do not meet the coverage rule. For the two-shock comparison the regions are grouped into four blocs: Mediterranean (48 regions), Continental (61), Eastern (28) and Nordic (15). The people-side of this

panel is education and health, since trust and participation are available only in survey rounds, and the full indicator list with Eurostat table codes is in Appendix B.

3.2 Methods

The analysis runs as a staged pipeline, and each stage is a separate test of the framework. On the wide panel the stages are principal-component pillar composites, a Gaussian graphical model of the indicator network, a consensus clustering with stability testing, a composite-level structural model and spatial diagnostics. On the annual panel they are relative-position series, a development-trap index and cross-lagged regressions. I use so many methods so that the framework is asked the same question in several independent ways, and the answer rests on whether they agree and does not depend on any one of them.

The composites summarise each pillar by the first principal component of its indicators, following the multivariate approach to composite indicators set out in the OECD and JRC handbook (OECD/JRC, 2008), computed in each imputed dataset and pooled across them. Section 4.1 shows that nothing depends on this weighting. The coupling between two pillars is the correlation between their composites across the 294 regions, estimated in each imputed dataset and combined on the Fisher z scale by Rubin’s rules. The scale of the coupling is examined by removing each region’s national mean and recomputing the correlation within countries.

The typology is built from the levels of the three composites (their correlations play no part), so that two regions fall in the same tier when their scores on Place-based Conditions, Human and Social Capital and Economic Activity are similar. In each of the twenty imputed datasets the regions are clustered by k -means on their three scores. The twenty partitions are then combined into one consensus partition by counting how often each pair of regions falls in the same cluster and clustering that agreement matrix (Strehl and Ghosh, 2002). I chose the number of tiers among four, five and six, with the stability of each cluster under bootstrap resampling of regions (Hennig, 2007) as the main criterion and the silhouette (Rousseeuw, 1987) as a secondary one. Four is the only number for which every cluster is stable by Hennig’s threshold of 0.75, with an average bootstrap Jaccard similarity of 0.87, and its silhouette is 0.389. The tiers are named from the data, by ranking their centres on the three composites from lowest to highest. The two calm windows are clustered separately. The 2014 to 2016 clusters are matched to the core clusters by the Hungarian method on the distance between their centres (Kuhn, 1955), so that a region’s tier can be compared between the two windows.

The indicator network is a Gaussian graphical model, in which each link between two indicators is their partial correlation with all other indicators held constant, estimated with the graphical lasso (Friedman, Hastie and Tibshirani, 2008) and selected by the extended Bayesian information criterion (Epskamp and Fried, 2018). Partial-correlation networks of development indicators are, to my knowledge, rare in regional studies. The method is well established in psychometrics (Epskamp and Fried, 2018; Haslbeck and Waldorp, 2020). Here the nodes are the development indicators themselves, which lets the syndromes be read straight off the network.

The remaining tools on the wide panel are standard, namely a confirmatory factor model, which compares one general factor with three pillar factors, and Moran’s I (Moran, 1950), which measures spatial clustering.

On the annual panel, the development 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 changes in relative position, in units of the 2008 cross-sectional standard deviation, and bloc gaps carry country-clustered bootstrap intervals. The pillar composites used for the annual correlations of Section 4.5 are equal-weight, sign-adjusted averages of the pillar’s indicators.

The development-trap index of Section 4.6 follows Diemer, Iammarino, Rodríguez-Pose and Storper (2022). Five-year growth in GDP per head, in labour productivity and in employment is compared against the European aggregate, the region’s own country and the region’s own preceding five years. A region is trapped when it fails more than half of the nine comparisons. Its construction is set out in Appendix B. The cross-lagged regressions of Section 4.8 regress each pillar on the previous year’s values of both pillars, in levels, with region effects and each country’s yearly mean removed. The coefficient on the other pillar therefore measures what it adds beyond the receiving pillar’s own persistence. Coefficients are in pooled standard deviations of each composite, with country-clustered bootstrap intervals from 500 draws.

Section 4: Results

4.1 Strong coupling in calm times

Before asking whether this structure constitutes an ecosystem, I need to establish that there is a strong, organised structure to account for in the first place, and the data show that there is a pronounced one. Unless stated otherwise, the figures in this section refer to the core period, 2017 to 2019.

Each pillar holds together as a coherent index, since its first principal component captures between 39 and 48 per cent of the variation within the pillar. These PCA-weighted composites correlate between 0.92 and 0.999 with a simple average of the same indicators, so the composites do not depend on how the indicators are weighted. The assignment of individual regions to tiers is more sensitive to how the composites are built, as Appendix A.9 shows.

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, as the table below shows. Correlations of this size would support combining the three pillars into a single composite.

Pillar pair	Before the pandemic, 2017 to 2019
Place and People	0.88
Place and Economy	0.89
People and Economy	0.85

Every pillar is strongly clustered in space, with a global Moran’s I between 0.84 and 0.90, all comfortably significant. The typology that emerges is geographically contiguous as well, with 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. Nobody disputes this, and it does not yet distinguish between readings of the structure. Coherent composites, close correlations and spatial clustering are all equally consistent with a single development gradient running across Europe. The calm-time picture does not reveal whether that gradient is best understood as one

dimension or as a coupled system of three. The system's behaviour under stress, the subject of Section 4.4, narrows that question without closing it, since it can show whether the pillars separate but not whether an imbalance between them is later corrected.


4.2 A reliable composite rather than three separable dimensions

The framework treats its three pillars as a way of organising the mechanisms and makes no claim that they are three separable dimensions of regional development, and the data bear that reading out.

At the level of individual indicators the three pillars cannot be told apart. A single general factor is preferred over a three-factor structure by a BIC margin of about 17, with or without the four anomalous indicators, 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. The comparison is about relative parsimony and says nothing about absolute fit, because both structures fit the indicator data loosely, as the confirmatory model in Appendix A.1 records. The point is therefore narrow: splitting the pillars into three factors adds parameters and returns nothing for them.

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

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




Nevertheless, the three-pillar composite achieves something no single pillar can. The typology it produces lies along a single development factor, with which it agrees at 0.925, 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 therefore more reliable than any of its components, even though those components are not statistically separable. Two indicators, overqualification and the belonging proxy, load against the people-side composite in every method (Appendix A.1).

However, a single factor does not account for everything. The map below shows, for each region in 2017 to 2019, its score on Human and Social Capital minus its score on Economic Activity, with both scores standardised across the 294 regions, so that a region at zero stands in the same relative position on both. A region in green ranks higher on its people-side (education, skills, trust and participation) than on its economy, which is the case across most of Spain, Greece, Sweden and

Finland. A region in red is the reverse, with an economy that ranks above its social indicators, as in Stuttgart, Oberbayern and Lombardy, and across much of Germany, Poland, Czechia, Slovakia and Hungary. The map shows where the two sides are out of step. A green region is not necessarily better off than a red one, since Andalucía is green and Stuttgart is red. The map cannot show whether this difference reflects a distinction between people and economy or only a single factor with uneven loadings, a question taken up in Sections 4.4 and 4.5, although it does show the configuration on which the shocks then act.

Human and Social Capital score minus Economic Activity score,
2017 to 2019, both standardised across the 294 regions



Green does not mean better off. The map shows where a region's social and economic sides are out of step, relative to the other regions.

A fairer alternative to one general factor is a general factor with a pillar-specific factor beside each pillar, since the pillars could share a common level and still each carry something of their own. I fitted that bifactor model on the same 294 regions and thirty-five indicators (mean of the twenty imputations, maximum likelihood). The general factor explains 76 per cent of the common variance, the economy-specific factor 10 per cent, the place-specific 7 per cent and the people-specific 7 per cent, with an omega hierarchical of 0.92. The three-factor model without a general factor is improper, with factor correlations at or above one. Whatever is specific to a pillar in the cross-section is therefore small, and smallest for people. For these reasons, the loosening between people and economy in Section 4.5 comes from a change in how the one common level is shared out between the pillars, since there is no hidden people dimension for it to come from.

For these reasons I treat the pillars as a way of organising the indicators, and the three together as a composite more reliable than its parts, without claiming that they are three separable dimensions of regional development. They are useful for what they produce, which is a stable typology, a layer of syndromes, and a separation of the pillars under the euro crisis.


4.3 The typology and the syndromes

The most concrete product of the framework is a four-tier typology of regions, built as described in Section 3.2, 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. I read the typology as a ladder because of this ordering on every pillar, which no single aggregate score would establish.

The ladder is stable across the two calm windows before the pandemic. Each tier attains a Hennig Jaccard bootstrap stability between 0.83 and 0.91 in 2017 to 2019, above the threshold of 0.75. The four-tier structure is already present in 2014 to 2016 across four imputers and two seeds, and from that window to the core one 86 per cent of regions remain in their tier. The high-performing tier retains 0.94 of its members, and upward moves outnumber downward ones about two to one (Appendices A.2 and A.5). Its geography is the one a regional economist would anticipate, with the high-performing regions in the north and west and the struggling regions in the south and east.

The four-tier typology, 2017 to 2019



Beneath the ladder lies a layer that a single composite index would not reveal, made up of four cross-pillar syndromes, each about 5.5 times stronger than a random grouping of the indicators, and it is visible only because the measurement spans three domains.

The typology is a cross-national development ladder observed at regional resolution. National context dominates the level, and how much it dominates depends on how finely the indicators are measured. Between-country differences account for between two thirds and nine tenths of the development gradient: 0.87 on the full indicator set, 0.81 with the ten indicators recorded only as national values removed, and 0.66 on the construct built only from regionally observed annual data, on the same regions. The full-set and clean constructs agree about where regions stand at 0.93 (0.96 within countries). The lower figure attaches to the more finely measured construct. All three sit where a panel spanning the European East-West gap would be expected to put them. However, the regional unit still adds information, since two thirds of the multi-region countries contain regions in more than one tier, several of them in three.

Section 4.4 follows the regions across two shocks, and Section 4.6 asks whether standing low on the ladder is the same condition as being stuck on it.

4.4 Two shocks, two signatures

The euro crisis moved Europe's regions along national lines, and fifteen years later they have not moved back, whereas the pandemic left them roughly where they were.


The wide panel describes the structure in a single calm window and says nothing about how long a region takes to recover from a shock, or whether 2020 resembled 2009. For that I use the annual panel of Section 3.1, in which the composite is fixed on the 2008 cross-section and each region is measured against the European average of its own year. This removes common growth and anything that struck the whole continent at once, so the figures below are changes in relative position.

Two episodes fall inside the window, and they behaved very differently.

In 2008 Mediterranean regions stood 0.42 standard deviations below the European mean. By 2013 they stood 0.94 below, a fall of 0.52, and 96 per cent of the 48 regions had lost ground. Over the same years the Continental core gained 0.28 and eastern regions 0.21. A gap of 0.80 standard deviations therefore opened between the Mediterranean and the Continental core in five years (country-clustered bootstrap interval 0.47 to 1.18).

The recovery is slow enough to be timed. The trough came in 2014, and in 2019, eleven years on, the Mediterranean was still 0.36 below where it had started. In 2023 it stood 0.64 below the European mean, against 0.42 in 2008. A decay fitted to the closing gap gives seven per cent a year, or a half-life of about ten years. However, that figure only describes the decay. The modelling companion decomposes the recovery and finds that it came from national shocks reversing and not from the regional system correcting, so the decay should not be read as a property of the regions.

Between 2019 and 2023, by contrast, the Mediterranean gained 0.14, while the Continental core lost 0.06, eastern regions 0.01 and Nordic regions 0.18. Only 23 per cent of Mediterranean regions lost ground in the pandemic, against 96 per cent in the crisis, a rate close to the 11 per cent of Continental regions that lost ground in the crisis. In the acute year, from 2019 to 2020, no bloc moved by more than 0.10, and the spread of regional changes was 0.14 against 0.45 in the crisis. The Mediterranean gain came afterwards, between 2020 and 2023. Over the whole period the South gained 0.32 on the Nordic regions, with a 90 per cent interval from 0.20 to 0.51 when countries are resampled within each bloc. The pandemic years therefore moved the blocs by well under half of what the crisis did, and they moved the South up. This is all the measure can show, since it cannot see a shock that hit every region equally and the pandemic's common damage lies outside it by construction.




152 NUTS 2 regions in 23 countries, fifteen indicators observed at regional level every year. Composite fixed on the 2008 cross-section; each region measured against the European mean of its own year, so common growth and Europe-wide shocks are removed.

Figure: the four bloc series in relative position, 2008 to 2023, with both episodes shaded. The Mediterranean falls to 2014 and recovers slowly, the east converges, and the pandemic years move the blocs little.

The finding survives changes to its own construction. Taking 2008 to 2010 as the base window gives a Mediterranean fall of 0.49. Weighting the fifteen indicators equally instead of by their first principal component gives 0.46. Counting Croatia as eastern changes nothing, and dropping Cyprus and Malta gives 0.54. Removing each country in turn moves the figure between 0.42 without Greece and 0.67 without Italy, so no single national experience carries it.

What fell Between 2008 and 2013 the Mediterranean's economic pillar fell 0.82 relative to Europe, while its human and social pillar fell 0.19 and its place-based pillar rose 0.10, so that almost all of the fall was in the economy.

The other two did not follow it down later. By 2019 the economy was still 0.65 below its 2008 position, whereas people and place had returned to where they started. The pandemic years show nothing of the kind in any pillar, with the economy up 0.13, people up 0.21 and place down 0.04.



Mediterranean regions (Greece, Spain, Portugal, Italy, Cyprus, Malta, Croatia), 48 of the 152. Each pillar is the mean of its standardised indicators, scaled by its own 2008 cross-sectional standard deviation.

Figure: the Mediterranean's three pillars in relative position, 2008 to 2023. The economy diverges from people and place through the crisis and is still below its 2008 position in 2023.

Part of this pattern is to be expected from what the pillars measure. The economy pillar is made of output and labour-market rates, which move with the cycle, whereas the people pillar is mostly educational attainment and the place pillar is life expectancy, road deaths and density, which are stocks that a recession does not run down. That the economy fell furthest is therefore unsurprising, and I do not rest the argument on it. The informative part is the persistence, since a cyclical fall in output and employment would ordinarily be made up within a few years, and this one was still there eleven years later.

Section 1.2 describes a region as a coupled system whose parts move together in calm conditions and separate under stress. The euro crisis shows this directly, on a shock large enough and slow enough to watch, and the separation lasted: eleven years after the crisis began, one part of the system was still displaced and the other two were not.

Which regions were hurt If some regions were fragile whatever the shock, the two episodes would have hurt the same places, and they did not, since of 152 regions none sits in the worst quarter of both episodes, where chance would put about ten. The correlation between how far a region fell in each episode is negative, at -0.48. Part of that is arithmetic, since a region already far below its starting point has more room to recover, and the same arithmetic bears on the overlap count. I therefore checked the count in three further ways. With the pandemic change taken net of how far each region still stood below its 2008 position in 2019, three regions are in the worst quarter of both episodes. With the pandemic measured from 2019 to 2021 only, six are, and within countries, with each country's mean change removed from both episodes, nine are, which is what chance would give. In every version the overlap is at or below chance. The regions hurt in the crisis were therefore no more likely than others to be hurt in the pandemic, and within countries the two episodes were unrelated.

Whatever made a region vulnerable in 2010 was therefore not what made it vulnerable in 2020, and any ranking of regions by resilience as a fixed property would have got the second episode wrong by relying on the first. I look for the difference in the shock and in the response to it, which Section 5.1 takes up.

Where the shocks entered The obvious candidate for the difference is what each region's economy is made of. The euro crisis followed a construction and credit boom, so one would expect regions with more construction employment to have fallen further. The pandemic closed borders and hotels, so regions living from tourism should have fallen further in 2020. I tested both with Eurostat's regional employment by sector and nights spent in tourist accommodation, taking the construction and public-employment shares of 2008 and tourist nights per resident in 2019 as the channels, each standardised across the 152 regions, with 90 per cent country-clustered bootstrap intervals.

Neither channel sorts regions within their country. Pooled across Europe, a standard deviation more construction employment in 2008 goes with a crisis fall of 0.15 (interval -0.24 to +0.07), but once country effects are included the relationship disappears (-0.01, interval -0.07 to +0.08). The public-employment share, the sheltered economy of Fratesi and Rodríguez-Pose (2016), goes with a smaller fall when pooled (+0.11) and with nothing within countries (+0.02). Tourism intensity does not sort the pandemic change in either direction (-0.00 pooled, +0.00 within countries). The same holds for the acute phase alone, where the coefficient within countries is -0.00 from 2019 to 2020 (interval -0.03 to +0.04) and -0.02 from 2019 to 2021 (-0.07 to +0.03). The result is therefore not produced by a contraction and a recovery cancelling over the longer window. Pooled, the two crisis channels account for 47 per cent of the gap that opened between the Mediterranean and the Continental core, but only because construction is higher in the South, and not because it sorted regions inside Spain or Italy.

The shocks did not enter through the regions' sectors: construction and tourism sort countries, not regions within them




Figure: the crisis fall against the 2008 construction share (left) and the pandemic change against 2019 tourist nights per resident (right), 152 regions, coloured by bloc. The slopes that appear pooled are between countries.


The incidence of the euro crisis therefore followed borders, and the obvious sectoral alternative does not hold. The pandemic, measured in the same way, entered through nothing that sorts regions within countries. The modelling companion reaches the same conclusion from the other direction, since its national crisis shock bears no relation to the regions' sectoral structure.

4.5 The people-to-economy link since 2012

The loosening between a region's people and its economy is a decade old: it began with the euro crisis recovery, the pandemic neither started nor interrupted it, and the link has not recovered since.

The wide panel shows the same fall across the pandemic, from 0.85 in 2017 to 2019 to 0.78 in 2021 to 2023. However, its post-pandemic window has measurement problems that this paper does not try to repair: the United Kingdom's regional data thinned after Brexit, one indicator exists only after 2021, and part of the panel changes resolution between the windows. I therefore rely on the annual series, which has none of these problems.

The annual panel gives the same correlation a longer run. Across the 152 regions it stands at 0.79 in 2008, rises through the crisis to 0.85 by 2011, holds there to 2013, and then declines without a break: 0.80 in 2016, 0.76 in 2019, 0.73 in 2021, 0.73 in 2023.



Correlation across the 152 regions in each year, pillar scores measured relative to the European mean. The decline begins with the euro crisis recovery and continues through the pandemic without a step.

Figure: the three cross-region pillar correlations by year on the annual panel, with the people-to-economy plateau at 0.85 in 2011 to 2013 and the decline since.

The rise to 2011 and the fall after it need to be read together with Section 4.4, since a correlation across regions and an imbalance within a region are different measures. In 2008 the Mediterranean's economy pillar stood 0.39 above its people pillar, each measured relative to Europe in its own 2008 standard deviation, which is the mark of the boom. The first years of the crisis removed that excess. By 2011 the gap was 0.03 and the two pillars ranked Europe's regions more alike than before, which is why the correlation rose. The economy then fell past the people pillar, to 0.24 below it in 2013 and 0.30 below it in 2019. The spread across the 152 regions of the gap between the two pillars follows the same path, falling from 0.64 in 2008 to 0.60 in 2011 and then rising to 0.73 in 2013 and 0.82 in 2021. That widening lies between countries, where the spread rose from 0.43 in 2011 to 0.70 in 2021, while within countries it stayed between 0.42 and 0.47 throughout. From 2011 onwards, the separation of the pillars in Section 4.4 and the decline of the correlation here are therefore one movement, and it takes place between national systems.

The series shows no pandemic step, so there was nothing to rebound from. The link did not recover after 2021 because it had already been sliding for ten years.

The result belongs with the decoupling literature of Section 2.2. Since about 2012 the tie between how a European region's people are faring and how its economy is faring has been loosening, and by 2023 the loosening had not stopped.

Dispersion across Europe and within countries The same panel shows where the crisis widened Europe and where it did not. Measured across all 152 regions, dispersion in relative development rose from 1.00 in 2008 to 1.19 in 2013 and has narrowed since, to 0.95 by 2023. Inside countries it barely moved, at 0.61 in 2008, 0.64 in 2013 and 0.61 in 2023, averaged over the seven countries with at least eight regions in the panel.

The crisis spread Europe's regions out along national lines and left the internal structure of countries as it was. This is the fact of Section 4.4 seen as a spread, and it is what a shock arriving through national fiscal positions would be expected to do.

4.6 Entrapment is a different condition


Being stuck and being poor are different conditions, and a typology of levels captures only the second.

Iammarino, Rodríguez-Pose and Storper (2019) describe European regions as settling into development groups with weak convergence between them, and Diemer, Iammarino, Rodríguez-Pose and Storper (2022) give the field a measure of being stuck. A region is in a development trap when its income, its productivity and its employment grow more slowly than the European average, more slowly than its own country, and more slowly than it did in its own recent past. These make nine comparisons, and a region is trapped when it fails more than half of them. I built that index for 236 regions in 27 countries from 2010 to 2023, with five-year windows, a wider set than the annual panel that includes Serbia and four member states the panel does not cover.

By bloc, the index agrees with Section 4.4. Mediterranean regions spent 9.3 of the fourteen years in the trap, a median of ten, with some of them trapped throughout, against 4.7 years for eastern regions. The Mediterranean share trapped ran between 0.88 and 0.94 from 2010 to 2015 and had fallen to 0.35 by 2023, while the eastern share fell to between 0.11 and 0.23 after 2017. One instrument measures relative position and the other measures growth against three benchmarks, yet they agree about which part of Europe lost a decade.

By tier, however, the index shows no gradient at all, with years spent trapped running at 7.2 for the struggling tier, 7.1 for catching-up, 6.8 for the stable core and 6.9 for high-performing.

Being stuck and being poor are different conditions



236 regions in 28 countries, 2010 to 2023. A region is trapped when income, productivity and employment each grow more slowly than the EU, than its own country, and than its own recent past, on more than half of the nine comparisons (Diemer et al., 2022), five-year windows. Roughly half the comparisons fail by construction, so differences between groups are interpretable and the level is not. Tiers from Section 4.3, 218 regions matched.

Figure: mean years trapped, 2010 to 2023, by bloc (a gradient from 9.3 to 4.7) and by tier (flat at 6.8 to 7.2).

Diemer et al. supply the reason when they report traps at low, middle and high income. The tiers of Section 4.3 record where a region stands, whereas the trap records how it is moving against three benchmarks. For instance, a prosperous region that has stopped growing is trapped, while a poor region that is converging is not. Nordic regions, second only to the Mediterranean at 8.2 years trapped, are a case of the first kind.

The two diagnoses should not be used interchangeably, and Section 4.7 shows they predict different things.

The index has one property that limits how it can be read. Each of the nine comparisons is made against an average or a previous period. Roughly half therefore fail by construction, and the median region sits near the threshold. Differences between groups and movements over time are therefore interpretable, but the absolute share of regions trapped is not.

4.7 Anti-establishment voting and the level of development

Within countries, anti-establishment voting tracks where a region stands, and how long it has been stuck adds nothing.

The typology and the vote share align along the whole ladder. 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. The gradient survives with country means removed, so it is not a country composition effect (Appendix A.10). 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. The obvious alternative reading is that voters respond to stagnation and not to position, and the trap index allows the two readings to be separated.

Across 213 regions in 25 countries the anti-establishment share correlates with the tier rank at -0.43 and with the share of the years from 2010 to 2019 spent in the trap at +0.38. Once national context is removed the two part company, since within countries the tier holds at -0.317 while the trap falls

to -0.009. Entered together with country fixed effects, the tier carries a standardised coefficient of -0.295 with a country-clustered standard error of 0.074, and the trap +0.014 with a standard error of 0.053. Both measures include years after some of the votes, so I repeated the comparison with each taken before the election. With the trap as the share of the years from 2010 to the year before the vote, the tier keeps -0.295 (0.073) and the trap +0.019 (0.055). With income per head averaged over the three years before the vote in place of the tier, income carries -0.257 (0.067) and the trap +0.011 (0.047).

The trap's raw association with discontent is a pattern between countries, because trapped regions sit disproportionately in poorer countries, which vote differently for their own reasons. Inside a country, the level predicts the vote and the trajectory adds nothing to it.

Two checks bear on how far this goes, and the first is a second election. At the national election before the last one, held between 2011 and 2017 depending on the country, the tiers order the vote in the same way, with anti-establishment parties at 36 and 31 per cent in the two lower tiers against 11 and 14 per cent in the two upper ones. Within countries, however, the gradient was weaker then, with a coefficient of -0.119 (0.065) on the tier, and the trap again adds nothing, at +0.025 (0.034). Measured before that vote, income carries -0.107 (0.047) and the trap +0.043 (0.029), which has the sign the trap argument implies and is not distinguishable from zero. The within-country gradient is therefore largely a development of the second half of the decade, and the change between the two elections shows it directly. Within countries the anti-establishment share rose more in the lower tiers, with a coefficient of -0.268 (0.102) on the tier, whereas years in the trap do not predict the rise, at -0.013 (0.053). One election further back, between 2006 and 2015, neither the tier nor income per head before the vote orders regions within countries, at 0.010 (0.083) and 0.032 (0.071), although across Europe the two lower tiers already gave more of their vote to these parties (32 and 24 per cent against 11 and 17). At that election the vote followed the six people-side indicators of the annual panel, measured before it, at -0.207 (0.069), an association that is absent at the last election, at 0.027 (0.092). The gradient in income is therefore recent, and an earlier one in education and inclusion had gone by the time it appeared.

The second check is a set of controls, on the 195 regions in 22 countries where they are observed regionally. Adding population density and net migration leaves the tier at -0.261 (0.087) and the trap at zero. Adding income per head, unemployment and tertiary attainment, which are among the things the tier is built from, takes the tier to -0.063 (0.059), with income per head carrying the association at -0.303 (0.107). On these controls the tier predicts the vote mostly because it orders regions by income. A fuller set changes that reading in part. With the share of the population aged 65 or over, the share of foreign citizens and the manufacturing share of employment added, on the 181 regions where all are observed, the tier stays at -0.261 (0.118) before income is entered and keeps -0.143 (0.055) beside income per head at -0.205 (0.061). Manufacturing regions give more of their vote to these parties at a given tier, at 0.115 (0.050), and it is with manufacturing held constant that the tier keeps an association of its own. Within countries, then, the vote follows a region's level, with income per head as the largest single part of it, while the trap stays at zero in every specification (Appendix A.10).

These remain associations across two elections. They do not overturn Rodríguez-Pose, Dijkstra and Poelman (2024), who relate the same index to the Eurosceptic vote with a different measure of the vote, a different horizon for the index and a different specification. The narrower claim they support is that, in this panel, discontent follows a region's level, chiefly its level of income, and not the length of time its growth has been stalled.

4.8 Do the pillars act on one another?

They do so in one direction only, and only within countries. Strong correlation between the composites is consistent with the pillars acting on one another, and equally consistent with a prosperous country lifting all three at once. The annual panel allows a direct check of whether a change in one pillar precedes a change in another once each region's own level and each country's common movement are removed.

Over one year, the effect of the people pillar on the economy is 0.013, with a 95 per cent interval from 0.022 below zero to 0.044 above, and it stays indistinguishable from zero however the people pillar is built. The reverse effect is small and positive within countries, at 0.130 with an interval from 0.041 to 0.209, and between 0.065 and 0.150 across four constructions of the people pillar, three of them with intervals that exclude zero. However, it does not survive when national movements are left in the data, where it turns negative. In other words, within a country a region whose economy improves tends to see its people pillar improve the following year, and the people pillar does not visibly lead the economy.

The pathway from people to economy that the framework asserts is neither identified nor refuted in these data, since a pathway running over longer horizons, or through channels these indicators do not capture, would look the same. The pathway from economy to people appears within countries only, and only at the one-year horizon reported here.


4.9 The pairwise literature within one structure

The conceptual companion began from a complaint, that the regional-development literature establishes the parts of the system two at a time and never assembles them (Section 2.3), and within this panel the assembled version can be checked directly.

For every pair of indicators that crosses a pillar boundary, the raw correlation across regions, the association a pairwise study starts from before its controls, is compared with the correlation that remains once a single general development factor has been removed. The first comparison asks whether the literature's relations appear in the data, and the second whether each is a direct connection or only 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. In other words, in this panel much of the association behind what the field has established one pair at a time lies on the same development dimension, which is the one-factor result of Section 4.2 again. A general factor built from the same indicators can also absorb a relation that is direct, so the attenuation shows how much of each association lies on the shared dimension and does not show which mechanism produces it. 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. 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.

Most pairwise findings are one shared dimension, a specific web survives



The relations that survive can be interpreted. The two findings from the literature that hold as direct connections, and are not reflections of the shared dimension, both concern institutions reaching into the social fabric, from government quality to civic participation and to institutional trust, the pathways of Charron and Rothstein (2018) and Cortinovis et al. (2017). By contrast, government quality to the economy and social capital to growth prove to be largely a reflection of the shared dimension. A caution is in order, since these are descriptive partial correlations that imply neither direction nor causation. The studies themselves use fixed effects and controls that a raw correlation lacks, so the comparison bears on the associations they start from and not on their estimates. The claim rests on the overall pattern and on no individual small relation, and the surviving web is cross-pillar, so 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 4.2.

Each study in the literature was therefore correct in what it found, although each observed a single relation and could not see how it fitted with the others. The ecosystem measurement recovers these relations, separates the shared development dimension from the specific connections, and presents them as one structure, which is the integrative contribution of the approach.

Section 5: Discussion

5.1 Why the Mediterranean lost a decade

Two large shocks left Europe's regions in different places, and the regions do not explain the difference. No region was among the worst hit in both episodes, so whatever made a region fragile in 2010 was

not what made it fragile in 2020.

The difference lies in how the shocks arrived and how they were met. The first came through national fiscal and financial positions (Lane, 2012), which is why its incidence follows borders so closely in Section 4.4 and why a region's own sectoral structure does not sort it. It was then met with consolidation in the countries it hit hardest, at a time when fiscal multipliers were larger than had been assumed (Blanchard and Leigh, 2013). The output lost to that consolidation was large (House, Proebsting and Tesar, 2020) and has been found to persist (Fatás and Summers, 2018). The second was met with job-retention schemes, with SURE (Council of the European Union, 2020), and with a common European fiscal response. At the regional level it entered through nothing that the sectoral data pick up, not even tourism. In my view, a shock differentiated by nation and amplified by national responses reorders the ladder, whereas a shock that is common and met in common leaves it in place. The countries of Appendix A.11 differ in how tightly their regions' social and economic conditions are tied, and the institutional differences behind these two responses would have to be traced there.

The data describe that contrast but cannot identify it, since there is one pair of episodes, no counterfactual, and no measure here of the policy response. The evidence supports three narrower statements: the damage followed national lines in the first case and did not in the second, the displacement persisted for more than a decade in the first case, and the regions hurt were different regions.

5.2 Level and trajectory

Sections 4.6 and 4.7 say something about which feature of a regional ecosystem turns into an outcome, since entrapment and the typology are unrelated and only a region's level predicts anti-establishment voting.

Section 1.2 describes regions as settling into configurations that are stable and able to shift. The two results add that political discontent corresponds to the level a region has settled at, which within countries is chiefly its level of income (Section 4.7), and not to the direction in which it happens to be moving. A region can stagnate from a comfortable position without that registering at the ballot box, and a region can be improving and still sit where people vote against the system.

I would argue that this favours measuring regional development as a level, which is what the typology does, since recent growth is a poor summary of a region's condition.


The companion papers use the word basin for such a configuration, and the evidence here fixes what the word can mean. The meaning rests on where the level is set. Between two thirds and nine tenths of the development gradient lies between countries (Section 4.3), and the modelling companion finds the same in its dynamics. In its account a region's level is pulled predominantly by its country, so the basin is national, and the regional part is the loop that sets how fast a region returns after an imbalance opens between its pillars. The modelling companion draws from this the implication that imbalances between the pillars arise at the level of the region, where cohesion money is also delivered, while the level of development is set by the country.

5.3 A system with a slow part

The place pillar varies less over time than the other two. In the wide panel, the place-based measures that move with development are the structural ones, namely regional vulnerability, government qual-

ity, broadband and research personnel, which correlate with the other two pillars at between 0.58 and 0.81. By contrast, the place-based measures that change from year to year, life expectancy, density, air quality and road deaths, correlate at between 0.14 and 0.51. On the 152-region annual panel, the share of each pillar's variance that lies within regions over time is 0.040 for Place, against 0.063 for People and 0.058 for Economy. Place therefore varies least over time relative to its spread across regions. A variance share of this kind is not a rate of adjustment, and the modelling companion estimates the persistence of each pillar directly. The JRC regional vulnerability series, which appears among the structural measures above, is modelled (Eklund et al., 2023) and shares inputs with the paper's own pillars, so it is not used here as evidence of speed.

Institutions can also be watched directly, through the European Quality of Government Index, which exists in five waves, 2010, 2013, 2017, 2021 and 2024 (Charron, Lapuente and Annoni, 2019, for the first four, and Charron, Bauhr and Lapuente, 2025, for the fifth). By tier its means barely move and the ordering never changes: the struggling tier sits at -1.3 in every wave until 2024, when it rises to -1.1, the catching-up tier at -0.4 to -0.5, the stable core at 0.5, and the high-performing tier at 1.1 to 1.3. On the 121 annual-panel regions where the index is published at NUTS 2, the share of its variance that lies within regions across the five waves is 0.055, which places it seventh of sixteen against the annual indicators at the same years. It sits above the stocks measured regionally, density (0.001), life expectancy (0.012), the low-education share (0.019), GDP per head (0.024) and tertiary attainment (0.032), and below every labour-market and education-flow indicator (0.06 to 0.19). On this measure, then, institutions vary as little over time as the stocks and less than every flow.



Thin lines are regions, thick lines tier means. Tiers are the four-cluster typology of the core window. NUTS 1 and national values are used where the index is not published at NUTS 2 (Charron, Lapuente and Annoni, 2019; QoG Institute, 2024 release).

Figure: the European Quality of Government Index by tier across five waves, thin lines regions, thick lines tier means.

In my view the place pillar is best read as a stock around which the people and economy pillars adjust. Under the euro crisis the economic pillar adjusted a great deal while the other two moved much less.

For the modelling paper this suggests an architecture in which the three parts of the ecosystem operate on different time scales, with the place-based foundation as the candidate slow variable. The recovery timed in Section 4.4 is a pattern for the companion to decompose, and it cannot be taken as a given rate. The modelling companion confirms the ordering but finds the margin small, with place only slightly more persistent than people, and it locates the slow part of the system in the level that the three pillars share. The pathways between pillars are estimated in the model, since Section 4.8 identifies only one of them from correlations alone, within countries and at one year.

5.4 What this implies for cohesion policy

Cohesion policy is usually judged by convergence, that is, by whether poorer regions close the gap on richer ones. On that test the period studied here is mixed, since eastern regions converged steadily while the Mediterranean moved the other way for five years and has recovered slowly since.

The two-shock comparison suggests a different test. The largest single movement in European regional standing over these sixteen years was a crisis-induced fall that one bloc had still not made up by 2023, and the second-largest shock produced almost no movement at all. If five years can open a gap that fifteen years have not closed, the value of cohesion instruments lies at least as much in preventing and shortening such episodes as in the slow work of convergence between them.

Two features of the structure bound what such instruments can reach. Most of the variation in regional development lies between countries, between two thirds and nine tenths of it depending on how finely the indicators are measured (Section 4.3). In addition, the part of a region's condition that correlates most strongly with the rest, its institutional and structural foundation, is the part that moves least from year to year. Regional instruments still have work to do, but the change they can bring about within a programming period is bounded, and the national layer does much of the work that regional policy is asked to account for.

Section 6: Conclusions

In this paper I have assembled an account of how the European regional economy is structured, and of how it behaved under two shocks, from several independent measurements. It rests on their agreement, since no single result is decisive.

Regional development is strongly and spatially organised (Section 4.1). The three pillars organise it as a concept and work together as a reliable composite, although they are not three separable dimensions (Section 4.2). The framework yields a four-tier ladder, stable across two calm windows, with a layer of syndromes beneath it (Section 4.3). This structure is organised predominantly at the national level. Between two thirds and nine tenths of the development gradient lies between countries, the lower figure attaching to the more finely measured construct, and the different constructs agree about where regions stand (Section 4.3).

The two shocks show the system behaving in two different ways (Section 4.4). The euro crisis reordered Europe's regions along national lines, pushing southern Europe 0.52 standard deviations down in five years and taking more than a decade to work off, whereas in the pandemic years no bloc moved by more than 0.18. Under the first shock the pillars came apart, with the economy taking almost all of the fall and staying down while the people pillar recovered, which is the separation the framework describes, and under the second nothing came apart. The regions hurt in one were no more likely than others to be hurt in the other, so these measures give no stable ranking of

vulnerability across the two episodes. Over the same years the tie between a region's people and its economy has been loosening, since about 2012 (Section 4.5).

Entrapment is a separate condition from a low standing on the ladder (Section 4.6), since a trap index built to the standard specification agrees with the two-shock account about which part of Europe lost a decade and bears no relation to the tiers. Within countries, anti-establishment voting follows where a region stands, chiefly its level of income, at the last election and more weakly at the one before, with controls and with each measured before the vote, and years in the trap add nothing (Section 4.7). Within countries, too, the economy leads the people pillar slightly from one year to the next, and the people pillar does not visibly lead the economy (Section 4.8). Beyond that, correlation between the composites is consistent with the pathways the framework asserts and equally consistent with a common national driver, and this panel cannot separate the two. Finally, the fragmented pairwise literature is recovered as one structure (Section 4.9), most of it the reflection of a single development dimension with a web of specific relations beneath.

The framework comes out of the test as the companion describes it, with one qualification. The system separates under a shock that enters through one pillar, settles into configurations that are stable and able to shift, and runs primarily through national systems with regional variation beneath them. However, the pathways the companion names between the pillars are, with one exception within countries, not identified in these data. My evidence therefore concerns what happens to the pillars under stress, and I do not observe them acting on one another. Whether the system is coupled in the stronger sense, with an imbalance between pillars later corrected, is the question the modelling companion estimates.

Three results become foundations for the modelling paper. The first is that development is organised through national systems that carry within-country regional structure. The second is that the place-based pillar varies least over time, which makes it the candidate slow stock of a dynamic account. The third is that the Mediterranean's recovery from the euro crisis, about seven per cent of the gap a year, is a measured decay, which the modelling companion decomposes and finds to be national shocks reversing, so that it is not a property of the regions. The model takes these as starting conditions, with regions nested within national systems and the place pillar as the candidate slow part, and asks what holds a region at its level and how fast it corrects an imbalance, which this paper can pose but cannot answer.

References

- Blanchard, O. J. and Leigh, D. (2013). Growth forecast errors and fiscal multipliers. *American Economic Review*, 103(3), 117-120.
- 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.
- Borsekova, K., Korony, S., Rodríguez-Pose, A., Styk, M. and Westlund, H. (2026). Decoding regional dynamics: institutions, innovation, and regional development in the EU. *Papers in Regional Science*, 105(1), 100133.
- 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., Bauhr, M. and Lapuente, V. (2025). Measuring the quality of government at the sub-national level: European Quality of Government Index (EQI), round 5, part I, technical report. Directorate-General for Regional and Urban Policy. Luxembourg: Publications Office of the European Union.
- Charron, N., Lapuente, V. and Annoni, P. (2019). Measuring quality of government in EU regions across space and time. *Papers in Regional Science*, 98(5), 1925 to 1954.
- Charron, N. and Rothstein, B. (2018). Regions of trust and distrust: how good institutions can foster social cohesion. In U. Bernitz, M. Mårtensson, L. Oxelheim and T. Persson (eds), *Bridging the Prosperity Gap in the EU: The Social Challenge Ahead*. Cheltenham: Edward Elgar, 220 to 242.
- Conway, A. (2020). Wellbeing and Macroeconomics: a SAGE approach. Master's thesis, Sciences Po Paris.
- Conway, A. (in preparation, a). A Regional Development Ecosystem Framework: Place, People and Economy in European Regions and Their National Systems.
- Conway, A. (in preparation, b). Slow Ecosystems: Regional Development, National Systems and the Limits of Cohesion Policy in Europe, 2000 to 2023.
- 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.
- Council of the European Union (2020). Council Regulation (EU) 2020/672 of 19 May 2020 on the establishment of a European instrument for temporary support to mitigate unemployment risks in an emergency (SURE) following the COVID-19 outbreak. *Official Journal of the European Union*, L 159, 1-7.
- Crescenzi, R., Luca, D. and Milio, S. (2016). The geography of the economic crisis in Europe: national macroeconomic conditions, regional structural factors and short-term economic performance. *Cambridge Journal of Regions, Economy and Society*, 9(1), 13 to 32.
- Diemer, A., Iammarino, S., Rodríguez-Pose, A. and Storper, M. (2022). The regional development trap in Europe. *Economic Geography*, 98(5), 487 to 509.
- Dijkstra, L., Garcilazo, E. and McCann, P. (2015). The effects of the global financial crisis on European regions and cities. *Journal of Economic Geography*, 15(5), 935 to 949.
- Dijkstra, L., Poelman, H. and Rodríguez-Pose, A. (2020). The geography of EU discontent. *Regional Studies*, 54(6), 737 to 753.
- Eklund, G., Sibilia, A., Salvi, A., Antofie, T.-E., Rodomonti, D., Salari, S., Poljansek, K., Marzi, S., Gyenes, Z. and Corbane, C. (2023). Towards a European wide vulnerability framework: a flexible approach for vulnerability assessment using composite indicators. JRC Technical Report, EUR 31515 EN. Luxembourg: Publications Office of the European Union.
- Epskamp, S. and Fried, E. I. (2018). A tutorial on regularized partial correlation networks. *Psychological Methods*, 23(4), 617 to 634.
- Fatás, A. and Summers, L. H. (2018). The permanent effects of fiscal consolidations. *Journal of International Economics*, 112, 238-250.

- Fingleton, B., Garretsen, H. and Martin, R. (2012). Recessionary shocks and regional employment: Evidence on the resilience of U.K. regions. *Journal of Regional Science*, 52(1), 109 to 133.
- Fratesi, U. and Perucca, G. (2018). Territorial capital and the resilience of European regions. *Annals of Regional Science*, 60(2), 241 to 264.
- Fratesi, U. and Rodríguez-Pose, A. (2016). The crisis and regional employment in Europe: What role for sheltered economies? *Cambridge Journal of Regions, Economy and Society*, 9(1), 33 to 57.
- Friedman, J., Hastie, T. and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3), 432 to 441.
- 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.
- Hennig, C. (2007). Cluster-wise assessment of cluster stability. *Computational Statistics and Data Analysis*, 52(1), 258 to 271.
- House, C. L., Proebsting, C. and Tesar, L. L. (2020). Austerity in the aftermath of the Great Recession. *Journal of Monetary Economics*, 115, 37-63.
- Iammarino, S., Rodríguez-Pose, A. and Storper, M. (2019). Regional inequality in Europe: evidence, theory and policy implications. *Journal of Economic Geography*, 19(2), 273 to 298.
- Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2), 83 to 97.
- Lane, P. R. (2012). The European sovereign debt crisis. *Journal of Economic Perspectives*, 26(3), 49-68.
- Lima de Miranda, K. and Snower, D. J. (2020). Recoupling economic and social prosperity. *Global Perspectives*, 1(1), 11867.
- Martin, R. (2012). Regional economic resilience, hysteresis and recessionary shocks. *Journal of Economic Geography*, 12(1), 1 to 32.
- Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, 37(1-2), 17 to 23.
- Muringani, J., Fitjar, R. D. and Rodríguez-Pose, A. (2021). Social capital and economic growth in the regions of Europe. *Environment and Planning A*, 53(6), 1412 to 1434.
- OECD/JRC (2008). *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing, Paris.
- Rodríguez-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.
- Rodríguez-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.
- Rodríguez-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.

- Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20, 53 to 65.
- Rubin, D. B. (1987). *Multiple Imputation for Nonresponse in Surveys*. John Wiley and Sons, New York.
- 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.
- Strehl, A. and Ghosh, J. (2002). Cluster ensembles: a knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3, 583 to 617.
- van Buuren, S. and Groothuis-Oudshoorn, K. (2011). mice: multivariate imputation by chained equations in R. *Journal of Statistical Software*, 45(3), 1 to 67.
- 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.