

Slow ecosystems: regional development, national systems and the limits of cohesion policy in Europe, 2000 to 2023

Conway

Contents

Abstract	1
1. Introduction	2
2. Where this sits, and what it validates	3
3. Data	4
4. The model	5
5. Estimation, checks and validation	7
6. How the regional ecosystem moves	11
7. Cohesion money	19
8. What cohesion policy can do inside this system	21
9. Discussion	23
10. Conclusion	25
References	26
Appendix A. Reproducibility	28
Appendix B. The nonlinear form and the feasibility domain	29

Abstract

Europe's regions have not converged in thirty years of cohesion policy, and the regions left behind are now where anti-system parties do best. Treating a region as an ecosystem of place, people and economy, I estimate its dynamics on 210 regions in 23 countries from 2000 to 2023 and test cohesion policy against them. Everywhere the system has one slow mode and two fast ones. The slow mode is a region's level, set mostly by the country, with a half-life of fifty-five to seventy years, and a century of regional income data agrees that less than half of a gap in level closes within a working life. The fast modes, imbalances between domains, are corrected within a generation. People lead the economy in the South, the economy leads people in the East, and in struggling regions the economy does not feed people at all. An economy hit alone, as in the euro crisis, keeps 52 per cent of its loss after twenty-five years in the South against 32 per cent in the East. Under maintained conditions about a seventh of regions stay in Europe's bottom quarter, seven of them under every assumption tried. Cohesion money, spread at an intensity that buys no relative gain, to regions that on the causal evidence absorb it least, in cycles far shorter than the decades the system takes to respond, could not have moved them. A policy fitting the system would be concentrated, sustained beyond a cycle and aimed at capacity where the economy cannot yet take the money.

1. Introduction

Regional inequality in Europe is a puzzle about persistence more than about growth rates. The map of lagging regions has barely changed since the Structural Funds were created. The regions that have not caught up are now the ones where anti-system parties win (Rodríguez-Pose, 2018). Iammarino, Rodríguez-Pose and Storper (2019) read this as clubs of regions on separate trajectories. Diemer, Iammarino, Rodríguez-Pose and Storper (2022) read it as a development trap that holds middle-income regions where they are. For Martin and Sunley (2015) it is a question of what lets a region resist, recover and renew. Each of these accounts describes the system through one of its symptoms, whereas in this paper I estimate the system directly.

This paper is the third of three. The first, Conway (2026a), proposed that a region is an ecosystem of three coupled domains: place-based conditions, human and social capital, and economic activity. Each supports the others, so that development is a property of the configuration and not of any one domain. It also named six directed pathways between them. The second, Conway (2026b), tested that proposal on European regions and found the configuration and its scale. It found one development dimension, set mostly at national level, and a euro crisis that moved regions along national lines and, within the South, pulled the economy away from people and place. It also found that the pandemic moved nothing, that entrapment and low standing are different conditions, and that from one year to the next there is only a small lead of the economy over people within countries. The second paper left the dynamics open, that is, how the domains act on one another, at what speed, and what that implies for a policy that has spent a generation trying to make regions converge.

In this paper I fit a dynamic model of the three domains to 210 regions over twenty-four years, with national conditions, cohesion payments and country shocks as drivers, and I test the estimator on simulated data before reading a single result. The model is the ecosystem's community matrix, the object ecologists estimate when they ask how a real community moves. It records how much each domain keeps of its state each year, how much it passes to the others, and what pulls it. Estimated by bloc, it gives the paper its central finding and its title. The regional ecosystem is slow, in that it corrects imbalances between its domains within a generation and changes its level little within a working life. The other results follow from this one, including why the crisis looked as it did, why the pandemic did not, where regions are heading, and why cohesion money could not have moved the regions it was meant for.

The paper makes four contributions, of which the first is the estimated dynamics of the regional ecosystem, by bloc, validated on simulated data and on the two shocks, with one slow mode and two fast ones and with the people-economy loop running in opposite directions in the East and the South. The second is a region-by-region decomposition of the euro crisis and the pandemic into national conditions, unexplained country shocks, cohesion money and the region's own dynamics. The third is a map of projected low standing under maintained conditions, with a core that holds under every assumption and agrees in direction with an independent trap index. The fourth is an account of why cohesion money, at the intensity it was spread and under the causal evidence on absorption, could not have moved the South, and of what would.

There are also several things the paper cannot do. The pathways are estimated as dynamic associations, and section 9.5 describes the design that would make them causes. I import the effect of a euro of cohesion money from the discontinuity evidence, because a panel of outcomes cannot

identify it, and I show why. Finally, a flexible learner given the same data forecasts no better than the ecosystem model does, which in my view says more about the data than about the limits of the model.

Section 2 places the paper, Section 3 describes the data and Section 4 sets out the model. Section 5 covers estimation, the recovery tests, identification, validation and a comparison with the shorter fit, and Section 6 reports the dynamics, the shocks and the projections. Section 7 turns to cohesion money and Section 8 runs the policy experiments, while Section 9 discusses the results and states the limits and Section 10 concludes.

2. Where this sits, and what it validates

Clubs, traps and places left behind. The club reading (Iammarino, Rodríguez-Pose and Storper, 2019), the trap index (Diemer et al., 2022) and the political geography of places that do not matter (Rodríguez-Pose, 2018) are three views of one thing, a set of regions that stay where they are whatever is done. I take them as evidence for the object I estimate and then use them to validate it. If the ecosystem is as the model says, slow along its level and fast across its imbalances, then clubs are positions along the slow direction, traps are regions whose level is pulled low, and the political map follows the level. Section 6.4 tests the last two, setting the model’s projected low-standing set against the Diemer index, which is built on a different logic, and against the empirical companion’s finding that anti-system voting tracks a region’s standing within its country.

Resilience. Martin (2012) and Martin and Sunley (2015) ask what governs a region’s resistance to and recovery from shocks. The model gives that question a number, the half-life of the fast modes, which is seven to twenty-four years depending on the bloc and the imbalance. It gives it in the terms ecology uses for the same question, since the model is the multivariate autoregressive community model that Ives, Dennis, Cottingham and Carpenter (2003) introduced for estimating a community’s stability from time series. I report the return rate and the reactivity they define in section 6.1 (Hampton et al., 2013, for the form with covariates). Dijkstra, Garcilazo and McCann (2015) and Crescenzi, Luca and Milio (2016) measured the 2008 crisis’s incidence, by settlement type and by national conditions. The empirical companion followed the displacement and added the pandemic, and here I decompose both.

Cohesion policy. The identified evidence comes from the eligibility discontinuity. Becker, Egger and von Ehrlich (2010) find that Objective 1 transfers raise growth in income per head by about 1.6 points a year within the programming period, with no employment effect. Pellegrini, Terribile, Tarola, Muccigrosso and Busillo (2013) find a positive effect on growth with the same design on regional data for 1994 to 2006. Becker, Egger and von Ehrlich (2013) find the effect confined to regions with enough human capital and good enough institutions, at about 1.1 points a year where the workforce’s upper-secondary share is above average and about 0.2, not significant, where it is below. Barone, David and de Blasio (2016) find Abruzzo’s gains undone when Objective 1 status ended. Becker, Egger and von Ehrlich (2018) find the same across Europe, where the growth gained under Objective 1 is more or less undone once a region loses the status. Two further results bear on where the money went and how far its effect is national. Rodríguez-Pose and Fratesi (2004) find that in Objective 1 regions the returns to infrastructure and business support, which took most of the commitments, were not significant, whereas education and human capital, about an eighth of commitments, had significant medium-term returns. Crescenzi and Giua (2020) find that the effect of cohesion policy differs by member state, with national conditions and models of intervention mattering beyond local conditions. I take the size and gradient of the effect from this evidence and

show in section 5.3 why a regional panel cannot supply them.

Ecosystem models and estimation. The community matrix is the object of Lotka-Volterra ecology (Bastolla et al., 2009; Rohr, Saavedra and Bascompte, 2014) and of its economic applications (Saavedra, Reed-Tsochas and Uzzi, 2009). Near the state a community occupies, any such model is well approximated by its linearisation, and the linearised form is what I estimate. The nonlinear form and its resilience measure are the paper’s interpretation and appear in Appendix B. The fitted model follows Doz, Giannone and Reichlin (2011), with a measurement model first and then a linear Gaussian state space estimated by the Kalman filter.

3. Data

Regions. The panel is 210 NUTS 2 regions in 23 countries, 2000 to 2023. It starts from the empirical companion’s regionally measured annual panel of 152 regions and adds 58 regions regained once that panel’s balanced-panel rule is dropped, above all in Germany, Poland, Greece and Austria. It was chosen so that every value is a measurement of the region it is attached to and none is carried down from a national aggregate. The United Kingdom is absent because Eurostat ceased publishing its regional series, Turkey, Norway and Switzerland fail the published panel’s coverage rule, and North Macedonia is dropped. A missing year is treated as a missing observation, which the Kalman filter skips. A pillar cell counts as observed when at least half of its indicators are present, in which case its composite is the mean of the indicators present divided by their mean loading. The share of pillar cells observed is 0.90 in 2000, 0.99 from 2009 and 0.978 overall. Each year’s European mean is computed over the 152 regions of the published panel, a fixed reference set that does not change with coverage. Fifteen indicators are observed annually. For the economy they are GDP per head in purchasing power standards, the employment, activity and unemployment rates, long-term unemployment and high-technology employment. For people they are tertiary attainment, the low-education share, early school leaving, adult learning, young people not in employment, education or training, and infant mortality. For place they are life expectancy, road deaths per million and population density, joined by three annual partial series, physicians per head, household broadband and motorway density, which replace the three-year window versions of broadband and motorways. Hospital beds were tested and excluded. Indicators where a higher value is worse are sign-reversed. The place pillar is the thinnest of the three at annual frequency, with three full indicators against six, and one of them, road deaths, moves with economic activity. Section 6.1 reports how place responds to national conditions with and without it, and section 6.5 tests the transition without it.

Additional measurements. The measurement layer also takes every further regionally measured indicator from the companion’s full set where it exists, as three-year window means (2014 to 2016, 2017 to 2019, 2021 to 2023). These are labour productivity, business and total R&D personnel, the at-risk-of-poverty rate, net migration, PM2.5 exposure and rail density. Five European Social Survey items on trust, contact, institutional trust, voluntary activity and civic participation enter in survey years (2012 to 2023) for the regions the survey covers. The European Quality of Government Index enters in its waves (2010, 2013, 2017, 2021), at NUTS 2 where it is published at that level and at NUTS 1 elsewhere. Ten indicators in the companion’s set are national values and are excluded, as is the Joint Research Centre’s regional vulnerability index, whose annual series is produced by regression filling and projection (Eklund et al., 2023) and so cannot serve as an observation. One further item, country attachment, failed a sign check specified in advance and was dropped.

The sparse measurements enter through two approximations. First, a three-year window mean is treated as an observation of the stock in the window’s middle year, although it is in fact the average

of three latent states. The recovery tests in section 5.2 simulated windows as true three-year means and estimated them this way, so the approximation has been tested. Second, a quality-of-government value published at NUTS 1 is assigned to each of its NUTS 2 regions as a separate observation with no shared error term, which overstates its information, and section 6.5 refits without those cells.

Cohesion money. The European Commission’s regionalised and modelled history of EU payments gives annual expenditure by NUTS 2 region, fund and programming period from 1989 to 2022, with a standard error on the modelled series. I use the European Regional Development Fund, the Cohesion Fund, the European Social Fund and the Youth Employment Initiative, as a share of regional GDP in current euros. I needed three code crosswalks, because the French regions carry their pre-2016 codes, five Greek regions switch to older codes for the 2014 to 2020 period, and five Polish regions renamed in the 2016 NUTS revision with unchanged boundaries are mapped to their payment codes. After them every region is matched to its payments. Croatia before accession has no payments and is set to zero with a flag. Belgian regional GDP for 2000 to 2002 is carried back from 2003 with a flag. The 2021 to 2027 period is not in the data, 2023 is absent, and 2022 is partial, so I carry the 2023 value from 2022 with a flag and never use it with a lag. Across 2000 to 2022 the median region received about 0.14 per cent of its GDP a year, the 95th percentile 3.3 per cent, and the maximum, in eastern Greece in 2013, 9.8 per cent.

National conditions. For each country and year from 2000 I take from Eurostat the ten-year sovereign yield spread over Germany and, for 2020 to 2022, the annual mean of the Oxford stringency index (Hale et al., 2021). Estonia, and the pre-accession and non-euro years without a benchmark yield, are set to zero with a flag, and for the non-euro members the spread also carries currency risk. The fiscal impulse is minus the annual change in the cyclically adjusted primary balance as a share of potential GDP, from the European Commission’s AMECO database (European Commission, 2026a), so that it measures discretionary policy and leaves out the automatic response of the budget to the cycle. Croatia’s first two years are missing and set to zero. EU pandemic money is the sum of Recovery and Resilience Facility grants and loans received in cash (Eurostat) and SURE loans, spread evenly over 2020 to 2022, per head.

Eligibility. Fifty-one of the 210 regions were below 75 per cent of the EU average GDP per head in 2008 (fifteen in Poland, six each in Bulgaria, Greece and Hungary, five in Czechia, four in Italy, three in Slovakia, two in Portugal, and one each in Estonia, Spain, Croatia and Latvia), and forty-one more were between 75 and 90 per cent.

4. The model

4.1 The ecosystem

For region i in country c , the model has three stocks $X = (P, H, E)$, place, people and the economy, each measured relative to the European mean of the year, and each year they evolve as follows:

$$X(i,t) = B X(i,t-1) + \alpha_i + G n(c,t) + \theta_0 \max(0, 1 + \beta z_H(i,t-1)) f(i,t-1) e_E + u(i,t)$$

B is the community matrix. Its diagonal gives how much of its state each stock keeps each year. Its off-diagonal entries are the six pathways of the conceptual companion ($B[i,j]$ is the effect of stock j on stock i). The term α_i is the region’s long-run intercept, a country part plus a regional part, and $n(c,t)$ holds the measured national conditions. The cohesion money share f is lagged a year, acts on the economy and is scaled by absorptive capacity. Here z_H is the region’s observed people score of the previous year standardised across Europe, with a missing year taking the nearest

observed value, θ_0 is the transfer's strength and β is how much it rises with human capital, floored at zero. Finally, $u(c,t)$ is a shock shared by every region in the country and $w(i,t)$ is a regional shock. The intercept is not the long-run level. With the inputs at zero the stocks settle at $(I - B)^{-1} \alpha_i$. With B close to a unit root the two differ by a large multiplier, which is why I work with finite-horizon projections (section 6.4).

Stochastic specification. The country part of α is $N(0, \Sigma_c)$ and the regional part $N(0, \Sigma_r)$, both diagonal, independent across countries and regions and carried as time-invariant states so that the filter integrates them out. The country shock $u(c,t)$ is $N(0, \Sigma_u)$, a full 3 by 3 matrix, independent over time, and the regional shock $w(i,t)$ is $N(0, \Sigma_e)$, also a full 3 by 3 matrix, independent over time and regions. The initial state $X(i,2000)$ is $N(0, V_0)$ with V_0 the first year's sample covariance of the composites, taken as independent of the intercepts, a restriction that the recovery tests share and do not test. Measurement errors are independent across indicators, regions and years. Since z_H is an observed score and not the latent people stock, the transition is linear and Gaussian given the data, and the Kalman filter's likelihood is exact for each country block. The exception is the 105 region-years, of 5,040, in which a people score before a region's first observation is filled from a later year. All fall before 2007, in 27 regions whose money share then averaged 0.12 per cent of GDP against 0.70 overall, and θ_0 and β are imported and not estimated, so the term they enter is a small fixed offset.

The ecosystem in this form. The matrix B has three modes, of which one is the level, with all three stocks moving together, and the other two are imbalances, with one stock above or below the others. The eigenvalues give each mode's half-life, and this is the language I use in section 6.

The measurement layer maps every indicator to its pillar's stock through a loading and an idiosyncratic variance. The annual indicators enter as the empirical companion's equal-weight pillar composites, divided by the pillar's mean loading so that the composite is unbiased for the stock, with a measurement noise variance estimated in the transition step. Where some of a pillar's indicators are missing, the composite is the mean of those present divided by their mean loading, which equals the published definition when all are present. The sparse indicators (window means, survey items, quality of government) enter through loadings estimated against those composites, corrected for the composites' transient noise. The layer relies on two approximations. A three-year window mean is treated as an observation of the middle year, which the recovery tests simulated as a true three-year mean and estimated this way. A quality-of-government value published at NUTS 1 is assigned to each of its NUTS 2 regions with no shared error term, and section 6.5 refits without those cells.

4.2 The nonlinear form and what it adds

Appendix B writes the same ecosystem as a generalised Lotka-Volterra system, in which each stock limits its own growth and the pathways enter a matrix A . Near a region's equilibrium its annual transition is $\exp(J)$, with J the Jacobian, and B above is the mean of $\exp(J_i)$ over regions. The nonlinear form adds a feasibility domain, the set of conditions under which all three stocks survive, which is ecology's measure of resilience. Calibrated to the fitted B and the companion's cross-section together, it misses both by about two standard deviations. I report the feasibility results in the appendix as the interpretation the nonlinear form gives, without treating them as results of the data. Everything in the body of the paper comes from the fitted B , except where section 7 says that a number comes from the calibrated population of Appendix B. In form, the state equation is the multivariate autoregressive community model ecologists fit to time series to estimate a community matrix and its stability (Ives et al., 2003; Hampton et al., 2013).

5. Estimation, checks and validation

5.1 The estimator

Estimation is in two steps (Doz, Giannone and Reichlin, 2011). The measurement step fits a one-factor model per pillar to the annual indicators, with uniquenesses bounded below to avoid improper solutions (Bollen, 1989), and loads the sparse indicators against the resulting composites. The transition step is a linear Gaussian state space per country, stacking the country's regions so that the shared shock and the country part of the intercept are carried in the covariance. It has two observation streams per region, year and pillar, namely the annual composite and the precision-weighted aggregate of whatever sparse indicators exist in that cell. Because positions are relative to a European mean computed from the 152 reference regions, which are part of the panel, the country blocks are not independent. The product of their likelihoods is a composite likelihood instead of the full joint one. However, the recovery tests below were run on data built and filtered the same way, and the estimator recovers the truth, which is the justification for using it. I checked the Kalman likelihood against the brute-force joint Gaussian on a small example (difference $1e-12$) and the smoother against the brute-force conditional expectation ($1e-14$). Priors are weak on B (diagonal centred at 0.8, off-diagonal at 0, standard deviation 0.2), on G and on the variances. Estimation is by maximum a posteriori, with uncertainty from a Laplace approximation at the mode, and every fit converged on the optimiser's own criterion.

5.2 Recovery tests

Before reading any result, I made the estimator recover known parameters from data simulated on the published panel's structure, meaning its 152 regions and their countries, the real national inputs and money, and the real measurement cells with their loadings and missingness. Transient measurement noise was sized from the data, and absorptive capacity was driven by the true people stock. I ran three tests, each with its pass mark fixed first at 0.05 per entry of B, a tolerance that is coarse for links of 0.005 to 0.025, so I rely on the finer statistics. With true values set away from the priors (off-diagonals of 0.01 to 0.08), the off-diagonal entries were recovered with root mean squared errors of 0.005 to 0.013 and the correct sign in 98 to 100 per cent of 50 replications. With the money effect fixed at the wrong value, twice the truth in size and double the absorptive gradient, errors were 0.006 to 0.015 and the sign was right in all 30. With true values set at the model's own estimates, which lie close to a unit root, off-diagonal errors were 0.002 to 0.006 against true values of 0.005 to 0.025, with the largest bias 0.003 and the sign right in all 16. In that test the spectral radius was recovered to the third decimal and the twenty-five-year cumulative responses to within two per cent. The long-run multipliers, the diagonal of $(I - B)$ inverse, were recovered with a median error of 13 per cent and a 90th percentile of 33, and for this reason I report destinations below as twenty-five-year projections instead of steady states.

I repeated the near-unit-root test on the long design before reading the long fit, with true values at the published pooled mode, data simulated on the 210 regions with their real missingness, sixteen replications and the money effect imported at the wrong value. The pass marks, fixed first, were an off-diagonal root mean squared error below 0.01, the correct sign in at least 90 per cent of entries, the spectral radius within 0.005 and the people-economy ratios within 0.12. Off-diagonal errors were 0.0015 to 0.0041 by entry, the sign was right in every case, the median error of the spectral radius was 0.0011 and the median worst error on a pathway ratio was 0.093. Eight of the sixteen replications met the optimiser's convergence criterion, and the recovered values are accurate in all sixteen. The

90 per cent intervals used below to call individual links distinguishable from zero contain the truth in 77.1 per cent of the 96 off-diagonal checks across the sixteen replications (56 to 94 per cent by entry, lowest for the links from the economy to place and to people). This falls just short of the band of 80 to 97 per cent set in advance, so the intervals are somewhat too narrow, and I read them with that figure in mind.

5.3 What the panel cannot identify

The money effect failed its recovery test, and the reason lies in the data and not in the estimator. Profiling the likelihood over θ_0 on data simulated on the published panel showed it flat. Values from near zero to one and a half times the truth lie within a third of a log-likelihood unit of one another, and the estimate lands wherever the prior puts it. Cohesion money goes to the same regions every year. In other words, 82 per cent of its variation is persistent difference between regions, which the region long-run levels absorb, and rightly so, since without them money is confounded with being poor and appears to cause decline. The remaining variation is year-to-year timing within regions, and at plausible effect sizes that signal is about two per cent of the noise. Eligibility changes do not help either, since sixteen regions of the published panel crossed the 75 per cent line between programming periods, eleven of them downwards because the euro crisis pushed them there. Every crossing changes funding only from 2021, which the payments data do not cover.

This is also why the identified evidence on cohesion policy comes from the discontinuity design. I therefore fix θ_0 at the Becker, Egger and von Ehrlich (2010) return of 1.2 euros of GDP per euro spent, mapped into the economy composite's units through the GDP indicator's loading (0.0077 composite units per percentage point of GDP a year). I set β at 0.87, which reproduces the 2013 contrast of 1.1 against 0.2 percentage points. All the money enters the economy stock, including the European Social Fund, whose programmes target skills and employment. The imported effect is an effect of the transfers as a whole on growth in GDP per head. No identified estimate separates what the Social Fund does to human capital, so I place the whole effect where the evidence measured it. The people stock may therefore gain from cohesion money in ways the model does not credit, and the absorption floor of section 7.2 is a statement about the growth effect only, and other returns to the money fall outside it. I report every money-dependent result at that value and across a grid of zero, half, one and two times θ_0 and zero, 0.87 and 1.74 for β . Refitting with θ_0 at zero and at twice its imported value moves no entry of B by more than 0.0017 and flips the sign of no national coefficient whose interval excludes zero.


5.4 Validation

I specified three tests in advance, in units of the development composite D, which is each pillar's composite divided by its 2008 cross-sectional standard deviation across the reference regions, averaged, and divided by the 2008 standard deviation of that average, relative to the mean of the reference regions each year.

Out of sample. I fit the model to 2000 to 2019 and forecast 2020 to 2023 from actual national inputs alone, with region intercepts at their filtered 2019 values and no outcome data after 2019.

root mean squared error of D	2020	2021	2022	2023
model	0.158	0.212	0.225	0.245
persistence (D held at 2019)	0.181	0.217	0.236	0.261
neural network country trend (2019 plus the country's mean 2015 to 2019 change)	0.158	0.216	0.231	0.260

The network is a multilayer perceptron with two hidden layers, trained on the same inputs, on complete rows only, and averaged over ten seeds. The model beats standing still in every year and the network from 2021, level with it in 2020, and it falls short of the country trend in 2023 (0.245 against 0.240). Figure 1 shows the four by year. The criterion fixed in advance required beating all three, and against the country trend the model misses it by 0.005. Relative regional positions over four years remain hard to forecast, and at that horizon this model adds nothing to a country's recent trend.



Forecasts use only the actual national inputs after 2019. The neural network is a two-layer perceptron trained on the same inputs, averaged over ten seeds. The ecosystem model beats persistence and the network in 2021, 2022, 2023, and in 2023 its error is 0.245 against 0.240 for a country-trend extrapolation.

Figure 1. Out-of-sample forecasts. Root mean squared error of D for 2020 to 2023 from a model fitted to 2000 to 2019, against a neural network, persistence and, for 2023, a country trend.

The euro crisis. The smoothed states reproduce the change in the Mediterranean minus Continental gap from 2008 to 2013, -0.50 against -0.49 in the data (95 per cent country-bootstrap interval -0.64

to -0.25), which is the check that the implementation is right. The empirical companion reports 0.80 for the same gap because it weights the fifteen indicators by their first principal component on 152 regions, whereas D averages the three pillars equally, and on those 152 regions D gives 0.54. The systematic path, driven by measured spreads and fiscal stance, long-run levels and money with no innovations, gives -0.14 (-0.15 to -0.12 across the money grid). In other words, measured national conditions account for about 28 per cent of the crisis gap. On that path the Mediterranean economy falls by 0.42 of its own standard deviation while people rise by 0.06, so the separation of the pillars under stress comes through the measured channel alone.

The pandemic. The systematic path from 2019 to 2023 moves no bloc's mean D by more than 0.15 (Mediterranean +0.08, Continental -0.05, Eastern +0.02, Nordic -0.08, against +0.09, -0.12, +0.05 and -0.02 in the data), so the model passes this test.

What the pathways add. I refitted the model with the six pathways held at zero, keeping the region and country intercepts, the national inputs and the country and regional shocks, both correlated across pillars. This is the account in which common national and regional factors and pillar-specific shocks do all the work. The log posterior falls by 82.9 for six parameters. Removing only the two links between people and the economy costs 39.4, and removing only the four links into and out of place costs 45.8. Without the pathways the slowest mode has a half-life of 33 years in place of 73, so the long persistence of the level belongs to the coupled system and not to any one pillar. The uncoupled model also forecasts 2023 less well, at 0.255 against 0.245. The correlation of regional shocks across pillars is itself estimated, at 0.21 between place and people and 0.27 between people and the economy. A model that takes those shocks as independent fits worse by 31.3 for three parameters and gives a slow half-life of 104 years, so the length of the slow mode is sensitive to this choice, although its order of magnitude is not.

5.5 The shorter fit

The paper's previous version was fitted on the empirical companion's balanced panel of 152 regions, 2008 to 2023, and I keep that fit as the comparison, on the fiscal measure it was estimated with, the unadjusted change in the primary balance. The first table sets the two pooled fits side by side, and the second the fits by bloc, with the half-lives of the slow and the two fast modes and the share of an economy-only shock left after ten and twenty-five years.

	2008 to 2023, 152 regions	2000 to 2023, 210 regions
persistence P, H, E	0.976, 0.960, 0.976	0.974, 0.972, 0.958
place persistence without road deaths	1.001	0.992
people to economy	0.78	0.63
economy to people	0.62	0.54
slow-mode half-life	130 years	73 years
slow-mode half-life across bootstraps	50 years to unbounded	44 to 389 years
fast-mode half-lives	24 and 12 years	23 and 12 years
slow direction (P, H, E)	(0.46, 0.57, 0.68)	(0.63, 0.56, 0.54)
reactivity	-0.009	-0.018
spread effect on the economy	0.059	0.052

bloc	window	half-lives (years)	economy left, 10 / 25 years	people to economy	economy to people
Mediterranean	2008 to 2023	238, 16, 16	0.90 / 0.80	1.21	0.55
Mediterranean	2000 to 2023	55, 21, 21	0.77 / 0.52	0.74	0.22
Eastern	2008 to 2023	144, 34, 8	0.55 / 0.30	0.32	1.04
Eastern	2000 to 2023	54, 25, 9	0.59 / 0.32	0.57	0.73
Continental and Nordic	2008 to 2023	102, 26, 6	0.61 / 0.42	0.91	0.70
Continental and Nordic	2000 to 2023	69, 23, 7	0.51 / 0.26	0.52	0.75

The architecture holds on both windows, with one slow mode and two fast ones in every bloc, a non-reactive system, the same national effect of the spread, and pooled people-economy ratios that lie inside the shorter fit's 90 per cent intervals. The Mediterranean looked more extreme in a window dominated by the euro crisis, whose displacement persisted inside it. With the pre-crisis years added its slow mode falls from 238 to 55 years and an economy-only shock leaves 52 per cent of the loss after twenty-five years, where the shorter fit gave 80. Economy persistence falls and people persistence rises, both outside the shorter fit's intervals. Over a window with a boom as well as a crisis the economy corrects faster and the people stock more slowly than the crisis window suggested. The Continental and Nordic loop changes direction (section 6.2), and Extremadura, which was in the shorter fit's low-standing core, leaves it (section 6.4).

6. How the regional ecosystem moves

6.1 One slow mode, two fast modes

Table 1 and Figure 2 set out the central result of the paper. In every bloc, the estimated community matrix has one mode that is close to a unit root and two that are not.

transition	slow mode: half-life, direction (P, H, E)	fast modes: half-lives, directions
all regions	73 years, (0.63, 0.56, 0.54)	23 years (0.95, -0.25, -0.21) and 12 years (0.51, 0.35, -0.78)
Mediterranean	55 years, (0.54, 0.50, 0.68)	21 and 21 years (0.96, -0.02, -0.03)
Eastern	54 years, (0.78, 0.52, 0.35)	25 years (0.99, -0.06, -0.16) and 9 years (0.57, 0.44, -0.69)
Continental and Nordic	69 years, (0.67, 0.56, 0.49)	23 years (0.73, -0.62, -0.29) and 7 years (-0.16, -0.48, 0.86)

The slow mode is the region's level, with all three domains moving together, and it moves slowly: its half-life is 73 years in the pooled fit and 54 to 69 years by bloc. The two fast modes are imbalances between the domains, one of them sitting above or below the other two, and the system corrects these within a generation. Which imbalance is corrected fastest depends on the bloc. In the Continental and Nordic regions it is people against the economy, with a half-life of seven years, and in the East

it is the economy against the other two, at nine years. In the Mediterranean both fast modes take twenty-one years, so its quickest correction is the slowest in Europe.

Figure 2 shows what this means for a shock. When the economy alone is hit, the Mediterranean system still carries 52 per cent of the loss after twenty-five years and the Eastern system 32 per cent. The loss therefore comes back faster in the East than in the South, and in both blocs the other two domains are slowly pulled down through the links. When all three domains are hit together, no bloc's system brings them back, because the shock has moved the level itself. In these data the regional ecosystem therefore keeps its domains in proportion and has little capacity to change its level.

Two kinds of shock: an economy hit on its own comes back faster in the East than in the South, and a level shift comes back slowly everywhere




Figure 2. Two kinds of shock. The response of the three pillars to a fall of one standard deviation in the economy alone (left) and in all three pillars together (right), from the transition fitted by bloc on 210 regions, 2000 to 2023.

The fast and slow modes are not pinned down equally well, although on the long panel both are finite in every resample. Across the twenty country-block bootstraps of the measurement panel, all of which converged, the second mode has a half-life between 17 and 30 years, whereas the half-life of the slow mode runs from 44 to 389 years (a spectral radius of 0.9844 to 0.9982). In other words, the data fix the generation, and in every resample the slow mode is finite and longer than forty years. Twenty resamples give a range and not a calibrated interval. At the bottom of that range a region still closes less than half of a gap in forty years. In the stability terms of Ives et al. (2003), the return rate is the dominant eigenvalue, whose half-lives by bloc are those of Table 1. The reactivity, the largest eigenvalue of $B'B$ less one, is negative in every bloc (-0.018 to -0.020) and pooled (-0.018), so no shock grows before it starts to decay. The diagonal of B is 0.974 [0.972, 0.977] for place, 0.972 [0.965, 0.979] for people and 0.958 [0.948, 0.968] for the economy. The people-economy ratios moved by at most 0.04 under every change to the measurement layer I tried (section 6.5).

The slow mode can also be checked outside the model. I measured how fast a region's relative income, its log GDP per head less the cross-regional mean of the same year, returns towards Europe's mean, as the slope of relative position on its value some years earlier across regions, annualised, with 90 per cent intervals from 500 country-cluster bootstraps. Over a century of the Rosés-Wolf regional database (version 6, Rosés and Wolf 2020; documented in Rosés and Wolf 2019), 173 regions in 16 countries at benchmark years from 1900 to 2015, the half-life is 54 years [43, 73], and from 1950 alone it is 42 [38, 53]. In ARDECO since 1980 (242 regions) it is 42 years [32, 71] for GDP per head and 52 [27, 131] for employment per head. Within countries the half-lives are longer, 61 years in both Rosés and Wolf windows and 79 to 186 years in ARDECO. The outside data therefore put the return

at forty to fifty-five years, about 1.3 to 1.6 per cent of the gap a year, and these are lower bounds, since measurement error biases the slope down. They measure a related but different object from the model's slow mode, the return to Europe's mean, where the model's is the return to a region's own level, and on both measures less than half of a gap closes within a working life.

Where the level comes from. The pull on a region's level is national. Over twenty-five years, 80 per cent of the variance of the cumulative intercept response is national for place, 97 per cent for people and 89 per cent for the economy. Where regions will stand in twenty-five years under maintained conditions is national by 47 to 72 per cent by pillar and by 62 per cent for the composite, since regional intercepts and shocks still show at a finite horizon. Where they stand now is national by 58 to 70 per cent, against the 0.66 the empirical companion reports on its regionally observed construct. The Lotka-Volterra form, calibrated separately to the companion's cross-section (Appendix B), puts the national share of intrinsic conditions at 0.93 (0.83 to 0.98). The two methods therefore agree that the level a region is pulled towards is set by its country. Crescenzi and Giua (2020) reach the same scale from the policy side, since what cohesion money achieves depends on the member state beyond local conditions.

Of the measured national conditions, two move the economy. A rise of one standard deviation in the sovereign spread lowers the economy stock by 0.052 [0.039, 0.065] in the same year. A fiscal expansion of one standard deviation, 1.5 points of potential GDP on the cyclically adjusted primary balance, raises it by 0.009, with a 90 per cent interval that just reaches zero (-0.0004 to 0.017), and by 0.011 [0.002, 0.020] when the structural primary balance, which also removes one-off items, is used from 2011. The measure matters for this result, since on the raw change in the primary balance the estimate is negative, at -0.007 [-0.015, 0.002], which is the sign the automatic deterioration of the budget in a recession gives. The fit is also worse by 1.4 log-posterior units, although no entry of B differs by more than 0.0002 between the two. The fiscal effect is about a sixth of the spread's per standard deviation. Consolidation after 2010 was also common to much of Europe. The Mediterranean's impulse relative to the European mean was large only in 2010, and the decomposition of section 6.3 attributes 0.002 of D to it in the Mediterranean by 2013, against 0.13 for the spread. Lockdown stringency shows nothing distinguishable from zero. EU pandemic money per head raises the people stock by 0.009 [-0.007, 0.026], which is not distinguishable from zero. The spread also goes with a slightly higher place composite (+0.018), and the fiscal impulse with a lower one, at -0.012 [-0.023, 0.001]. Road deaths, which fall with economic activity, are the obvious suspect for both, but they are not the explanation. A refit without them leaves the two point estimates close to where they were, at +0.023 and -0.013, and only widens their intervals to include zero, while the fiscal effect on the economy is unchanged at 0.009 [0.000, 0.017]. Place is close to a unit root in that construction (section 6.5), so I do not interpret either coefficient.

6.2 The pathways, and which way the loop runs

The cross-links in B are small year to year (0.001 to 0.026), which is why the empirical companion, working from correlations, found only one of them. However, against the decay rates implied by a diagonal of 0.958 to 0.974 they compound. Figure 3 shows what a permanent programme does, by bloc: the funded domain rises for decades and the other domain follows it through the links. As a share of the funded domain's own response, the economy reaches 7.6 per cent of a people programme's effect after one seven-year programming period and 26 per cent after twenty-five years. People reach 4.6 and 17 per cent of an economy programme's, with more where the loop is strong. In the ecosystem model's units, where a pathway is the annual link divided by the receiving stock's decay, people to economy is 0.63 [0.52, 0.75] and economy to people 0.54 [0.33, 0.72] pooled. These values are stable

across the money grid, and the recovery test on the long design returned the ratios with a median worst error of 0.093. Twenty country-block bootstraps of the measurement panel give 0.42 to 0.89 for people to economy and 0.37 to 0.66 for economy to people, and the 90 per cent intervals covered the truth in 77.1 per cent of simulated checks.

Two permanent programmes: the pillar you fund rises for decades, and the other pillar follows it, slowly




Figure 3. Two permanent programmes. The cumulative response of people and the economy to a rise of a tenth of a standard deviation a year in the conditions of one pillar, by bloc.

The loop runs in different directions in different parts of Europe. Letting B differ by bloc improves the fit by 33.7 log-posterior units for eighteen parameters and splits the people-economy loop. In the Mediterranean people lead the economy (0.74 against 0.22). After twenty-five years the economy has reached 21 per cent of a people programme's effect while people have reached 5 per cent of an economy programme's. In the East the economy leads people (0.73 against 0.57). On the long panel the economy also leads people in the Continental and Nordic regions (0.75 against 0.52), which is a change from the shorter fit, where people led the economy in that bloc (0.91 against 0.70). The Eastern pattern looks like catching up through investment and employment, whereas the Mediterranean pattern is the one the conceptual companion proposed, and it is now the only bloc where people lead. Pooled numbers are therefore an average of these systems. Grouping countries by welfare regime in place of geography fits better still, and section 9.4 reports it. Fitting by the empirical companion's tiers instead of by bloc adds one result. In the struggling tier the economy does not feed people, and the estimated link is negative (-0.65), while people feed the economy strongly (0.86). People also lead in the catching-up tier (0.44 against 0.18), and the economy leads in the stable and high-performing tier (0.71 against 0.49). The regions the companion calls struggling are the ones whose ecosystem has lost the economy-to-people link.

One mechanism for the Eastern direction can be checked on observed data, namely that a growing economy draws migrants and migrants raise the people stock. On the observed composites with region effects and country-year means removed, the Eastern economy at one year does not raise net migration the next (-0.03, interval -0.28 to +0.08), although net migration still raises the people composite the year after (+0.06, interval +0.02 to +0.08). Migration is therefore not the mechanism of the Eastern direction, and no bloc meets the mediation rule I set for it.

Two readings of the pathways are open. Under the conceptual companion's reading they are the mechanism of a slow ecosystem, whereas under the sceptical one they are dynamic associations produced by whatever moves a region's domains together over decades, and the estimator would find them either way. These data cannot separate the two readings, so "pathway" in this paper names an

estimated association, and section 9.5 states the design that would decide between them. The speed of the system, by contrast, holds under either reading, and it is what the rest of the paper builds on.

6.3 The two shocks, region by region

The transition is linear, so each region's smoothed path decomposes exactly (Figure 4) into the contributions of its start and long-run level, each national condition, cohesion money, the country shocks and its own shocks. Over 2008 to 2023 the sovereign spread cost Mediterranean regions 0.21 of D and Eastern regions 0.13, and gave Continental and Nordic regions 0.15, the mirror image of a Europe-relative measure. EU pandemic money gave the Mediterranean 0.07. Cohesion money, at the imported effect, gave the East 0.06 and the Mediterranean 0.01.

What moved Europe's regions: the crisis came through national conditions and national shocks, and the recovery was those shocks reversing




Figure 4. What moved Europe's regions. Each bloc's mean change in D split into the sovereign spread, unexplained country shocks, EU pandemic money, cohesion money at the imported effect, and the start and long-run pull, for 2008 to 2013 and 2008 to 2023.

The euro crisis, taken as the change from 2008 to 2014 against a counterfactual that holds 2008 sovereign and fiscal conditions and removes country shocks, cost Mediterranean regions 0.35 of D . Of that, measured conditions account for 0.07 and country shocks for 0.21 (the two counterfactuals are computed separately and do not sum exactly). The 28 per cent of section 5.4 is a different object, namely the measured share of the 2013 gap against the Continental bloc on the systematic path. The Continental bloc gained 0.25 and the Nordic 0.11, while the Eastern bloc lost 0.06. The country shock is the part of each country's common movement that the measured conditions do not explain, in other words a residual. It establishes that the crisis loss was national, common to every region in a country, and not captured by bond markets or by fiscal stance, whether adjusted for the cycle or not (section 6.1), but it does not establish what the loss was. However, it was not sectoral incidence in disguise. Regressing each region's cumulative shock contribution on its 2008 construction share (Eurostat nama_10r_3empers) gives -0.07 $[-0.11, +0.03]$ per standard deviation pooled and nothing within countries. Across the 23 countries the Spearman correlation between a country's construction share and its shock is -0.12 . The country shock was therefore national, and it was not the sum of the regions' building booms. The recovery is more informative, since the country shocks' contribution to Mediterranean positions goes from -0.22 in 2013 to -0.05 by 2019, where it stays in 2023, while the spread's contribution persists and grows, from -0.13 in 2013 to -0.20 by 2019. In other words, southern regions regained ground mostly as the shocks that had pushed them down reversed.

I also asked whether the housing and credit cycle accounts for the country shock, by entering one further national condition at a time in place of lockdown stringency (Eurostat, 2026). House-price growth is associated with the economy almost as strongly as the spread, at 0.048 [0.040, 0.057] per standard deviation. It improves the fit by 42.2 log-posterior units and lowers the variance of the economy's country shock by 17 per cent over the whole period. Private credit flow has a smaller effect, at 0.016 [0.008, 0.024], the current account a negative one, at -0.018 [-0.029, -0.006], and the change in private debt none that is distinguishable from zero. In the crisis decomposition, however, house prices take their share from the spread. They account for 0.05 of the Mediterranean's position in 2013, the spread's contribution falls from 0.13 to 0.08, and the contribution of the country shocks barely moves, from 0.22 to 0.21. The housing cycle and the sovereign spread therefore appear to measure the same national financial conditions, and neither accounts for the residual.

Two experiments bear on the pandemic, and they carry different weight (Figure 5). The first is a counterfactual within the model, in which removing RRF and SURE money, the one measured input that moved in 2020 to 2023, leaves the Mediterranean at +0.02 instead of +0.09. Because European money went where the pandemic hit hardest, its effect and the shock's own incidence cannot be separated with these data, and the money's estimated effect on the people stock, 0.009 [-0.007, 0.026], is not distinguishable from zero. The second is a stress test, and should not be read as a counterfactual, in which replaying each country's 2010 to 2013 deviations in spread, fiscal stance and estimated country shock onto 2020 to 2023 turns the Mediterranean's gain into a loss of 0.20. Since the replayed country shock is a residual from a different crisis, this shows what a shock of that size and national incidence would do to the fitted system and says nothing about why the pandemic did not produce one. For these reasons the two should be read separately and not combined into a claim about the pandemic's shape.




Figure 5. The two shocks in the model. The euro crisis against a no-crisis path, 2008 to 2014 (left), and the pandemic, 2019 to 2023, as it happened, without RRF and SURE money, and under a stress test (right).

6.4 Projected low standing under maintained conditions

A projection is a region's position twenty-five years on from its 2023 state, with national conditions held fixed and its 2014 to 2020 average cohesion money. Absorptive capacity follows the region's projected people stock (no observed score exists for future years), and the transition is fitted by bloc. The recovery test showed the long-run multipliers to be fragile at these persistence levels and the


twenty-five-year responses to be estimated within two per cent, so projections are the object the data support. A region has projected low standing if its projection lies below the 25th percentile of the 2008 distribution. This is a scenario classification, showing where the ecosystem carries a region if the conditions named are maintained. A model that forecasts four years ahead no better than a country trend cannot say which conditions will hold, and for this reason I give the classification across conditions, transitions and thresholds.

Under 2019 conditions, 13.8 per cent of regions (29) have projected low standing, about a seventh, made up of 40 per cent of Mediterranean regions, 12 per cent of Eastern and none of the Continental or Nordic. The set depends on the conditions held and on the threshold. Across thresholds at the 20th, 25th and 30th percentiles and the scenarios for conditions the share runs from 7.1 to 24.3 per cent. Figure 6 maps the baseline (bloc transition, 2019 conditions, 25th percentile of 2008) as a fate, showing whether a region is below that line today and whether it is below it in 2048. Twenty-eight stay below it, eighteen climb out, one falls in, and 163 stay above, so 29 are below the line in 2048. The one that falls in is Piedmont. The eighteen that climb out are four Polish, three Hungarian, three Spanish (Extremadura, Castilla-La Mancha and Andalusia) and two Bulgarian regions, with Hainaut, Champagne-Ardenne, Corsica, the Azores, Madeira and Stredné Slovensko. Seven regions are below threshold under every combination of transition, conditions and threshold, and are hatched on the map: Severozapaden in Bulgaria, the Spanish autonomous cities of Ceuta and Melilla, Campania, Apulia, Calabria and Sicily. Twelve are below the 25th percentile under every condition and transition, adding Severen tsentralen in Bulgaria, Eastern Macedonia and Thrace, Northern Hungary, Molise and Východné Slovensko. Under 2023 spreads and fiscal stance, with pandemic inputs at their pre-2020 levels, the share is 16.7 per cent and the set overlaps the baseline at 0.75, with one Greek region leaving and Hainaut, Yuzhen tsentralen, Severozapad in Czechia, three Hungarian and two Polish regions entering. Across the twenty measurement bootstraps, recomputed with each draw's transition and the other parameters at the mode, five of the seven core regions are in low standing in every draw, Ceuta in seventeen of twenty and Melilla in fifteen. Of the 28 regions in low standing under the pooled transition, 26 hold in at least 90 per cent of draws, the two exceptions being Melilla (75 per cent) and Ceuta (85). Sixty-two, 30 per cent, are below threshold under at least one combination, and my claim concerns the core and the range.

Who stays in the bottom quarter, who leaves it, who falls into it

Below the 2008 bottom-quarter line, in 2023 and in projected 2048

- stays in the bottom quarter (28)
- falls into the bottom quarter (1)
- climbs out of the bottom quarter (18)
- stays above (163)
- ▨ stays there under every transition, condition and threshold (7)



Each region is placed by whether its composite position is below the 2008 bottom-quarter line today (2023) and in 2048, projected under maintained 2019 national conditions and 2014 to 2020 cohesion money with the transition fitted for its bloc. Hatched: in the bottom quarter in 2048 under every combination of transition, conditions and threshold. Pale grey: not in the panel.

Figure 6. Projected low standing. Regions below the 2008 bottom-quarter line in 2023 and in 2048 under maintained 2019 conditions, with the seven that stay below it under every combination of transition, conditions and threshold hatched.

At this point the club and trap literatures can be used to check the model. The Diemer et al. (2022) trap index is built from relative growth and mine from projected level, so the two should agree if both see the same system. I compare each region's projected level with its years in the trap, for the regions the index covers, which do not include Belgium. The two agree in direction, but more weakly than a common system alone would imply, with a Spearman correlation of -0.20, and the share with projected low standing rises from 8 to 15 to 23 per cent across the index's terciles. Finally, the empirical companion's finding that within countries anti-system voting tracks where a region stands is consistent with reading the political map of Rodríguez-Pose (2018) as a map of the slow direction, since those regions stand where the ecosystem is pulling them, although the model contains no vote data and does not test this.

6.5 Sensitivity of the transition to the measurement layer

Four refits of the full model, each from the converged mode, remove one part of the measurement layer at a time, namely the sparse indicators, the NUTS 1 quality-of-government cells, the survey items and road deaths. None of them changes the people-economy ratios by more than 0.04. Three of them leave the low-standing set almost unchanged (overlap 0.97 to 1.00 with the baseline). The fourth, without road deaths, overlaps it at 0.68, the same sensitivity the shorter fit showed.

The refit without road deaths is the one that exposed the place pillar on the shorter fit, where life expectancy and population density alone, both of which drift within regions and hardly revert, took the place stock's persistence to 1.001, a unit root. On the long panel place is stationary in every construction, at 0.974 with road deaths and 0.992 without. The three annual partial series of section 3 enter here. Physicians and household broadband end in 2021 in Eurostat's regional tables and motorway kilometres are missing for Germany and Greece, so none met the published panel's admission rule of near-complete coverage, but the Kalman filter takes them as partial series. Motorway density has no coverage in 2008 and could not be tested against the composite in that year. I therefore admitted it on its first well-covered year, where it correlates 0.68 with place in 2013, and it passes the build's own sign check (0.64 in 2019). Hospital beds were tested and stay out, as do tourist bed places, which do not correlate with the composite, and R&D personnel, which is published at NUTS 1 for France, Belgium and the Netherlands.

With place stationary, the place-side pathways are estimated. Economy to place is 0.78 [0.60, 0.95], place to economy 0.10 [0.01, 0.20], place to people 0.13 [-0.01, 0.28] and people to place -0.04 [-0.20, 0.11]. The economy therefore builds place and place feeds the economy weakly, while its effect on people and the effect of people on place are not distinguishable from zero. Without road deaths the economy-to-place ratio divides by a decay of 0.008, and I do not interpret it, although the annual link it rests on is 0.020 to 0.029 across constructions. Place remains the most persistent of the three stocks (0.974 against 0.972 for people and 0.958 for the economy), though its interval overlaps that of people, so the slow part of the system is the common level of section 6.1 more than any one stock, while the people-economy loop is the same on every measurement layer I tried.

7. Cohesion money

7.1 Relative intensity

Cohesion money's contribution is Europe-relative, like everything else in the model. The money term enters the state equation as a share of GDP and is not centred there. Positions are reported in D, which subtracts the mean over the reference regions in each year, so a region's reported gain from money is its effective intensity, the share scaled by its absorptive capacity and carried through B, less the reference mean of the same quantity. This is the measure the Treaty sets, since Article 174 of the Treaty on the Functioning of the European Union gives cohesion policy the aim of reducing disparities between the levels of development of the regions and the backwardness of the least favoured ones (European Union, 2016). At the imported effect, it lifts the destination of eastern regions by 0.086 of D [0.075, 0.100] (zero to 0.22 across the money grid) and of Mediterranean regions by 0.011 [0.009, 0.012] (zero to 0.04). The difference comes from the data on where the money went, not from the model, since eastern regions received well above the European average per unit of GDP and southern regions received about the average. Money at the average effective intensity therefore produces no relative gain. Set against the crisis, cohesion money gave the Mediterranean 0.01 of D over the whole

decomposition window, whereas the crisis cost it 0.35 between 2008 and 2014. In relative terms there was almost no cohesion advantage for the crisis to erase.

7.2 Absorptive capacity

The absorptive term is imported, with θ_0 from Becker, Egger and von Ehrlich (2010) and β of 0.87 from the 2013 contrast (section 5.3), and in the fitted ecosystem its consequence is a floor (Figure 7). With people scores standardised across Europe, an eligible region whose score is low enough that $1 + \beta z_H$ falls to zero absorbs nothing. Under the people-stock measure fourteen of the fifty-one eligible regions are on that floor (five of Bulgaria's six regions, all but Yugozapaden, which contains Sofia, together with Campania, Apulia, Calabria, Sicily, Eastern Macedonia and Thrace, the North Aegean, Northern Hungary, the Azores and Východné Slovensko), and under the upper-secondary measure nine are. A seven-year transfer of one percentage point of GDP raises the economy pillar of an eligible region off the floor by 0.033 of its standard deviation by year seven, a near-eligible region by 0.025, and a region on the floor by nothing. An eligible region off the floor therefore absorbs more than a near-eligible one. The line that separates regions that absorb from regions that do not is the floor and not the eligibility threshold, which is how the Becker et al. (2013) gradient appears on the map.

The absorption floor: on the causal evidence, the regions cohesion money is for are the least able to use it




Figure 7. The absorption floor. The share of a transfer's effect that each eligible and near-eligible region absorbs under the gradient of Becker, Egger and von Ehrlich (2013), with capacity measured by the people pillar (left) and by the upper-secondary share (right).

Whether the floor exists depends on two choices that the evidence does not make for me, and I report the results under both. The first is the capacity measure, because Becker et al.'s moderator is the workforce's upper-secondary share, whereas the people stock is the framework's broader human and social capital, which loads on tertiary attainment, adult learning, early leaving and youth inactivity. The two sort the eligible regions in opposite ways, with ten of the fifty-one above the European mean on the people stock and thirty-seven on the upper-secondary share, because eastern and southern regions combine high secondary completion with low tertiary attainment. The second is which moments to match, since Becker et al.'s below-average effect is not distinguishable from zero. Appendix B fits the term on the calibrated population under each choice, and the floor ranges from none to three quarters of eligible regions across defensible choices. The truncated linear form is a third choice, and a smoother response would soften the floor without removing the gradient. The gradient itself

survives every choice, whereas the comparison across the eligibility line does not, since near-eligible regions absorb slightly more than eligible ones under the stock channel and slightly less under the flow channel (Appendix B). Capacity is also not a fixed attribute of a region, since Rodríguez-Pose (2026) finds cohesion spending associated with faster growth only where the quality of government improved.

7.3 The channel

The fitted persistence says something about how cohesion money worked when it worked. In the model, a transfer raises the economy's conditions, the stock adjusts slowly toward a higher level, and when the transfer stops the stock decays slowly. On the fitted transition, 63 to 104 per cent of a seven-year push is still in place fourteen years later (section 8). Barone, David and de Blasio (2016) found the opposite in Abruzzo, where the gains were undone when Objective 1 status ended, and Becker, Egger and von Ehrlich (2018) found the same across Europe. The model and this evidence would be consistent if cohesion money acted as spending, a flow that raised measured activity while it flowed and vanished with it, and not as a change in what the region could do. The calibrated population of Appendix B can carry either channel, and I ran both. Both fit the Becker moments exactly, since a seven-year change constrains size and gradient but not persistence. They differ only in what remains when the money stops, which is nothing under the flow channel by construction and most of it under the stock channel. The discontinuity evidence therefore cannot choose between them. The evidence on what happens when status ends, the only evidence that can, favours the flow, in Abruzzo and across Europe. This is a comparison of two channels against two studies of status loss, which falls short of a test. If cohesion money was a flow, that is also why a model of regional stocks cannot see it in the data.

8. What cohesion policy can do inside this system

Every experiment in this section runs on the fitted ecosystem, by bloc, from each region's 2023 state under 2019 conditions. An intervention is a rise of a tenth of a pillar's 2008 standard deviation in that pillar's annual inflow of conditions, and responses are in D. None of these experiments has a pass mark, and every result is given across the money grid and, where it matters, the two capacity measures.

How long. A sustained push realises a quarter to three eighths of its twenty-five-year effect (27 to 37 per cent) by the end of one programming period, in every bloc and for every pillar. Seven years are 28 per cent of twenty-five, so on the fitted persistence the effect accumulates almost in proportion to the time the push is held. A push held for seven years and then removed leaves 63 to 104 per cent of its value in place fourteen years later (above one where the other domains keep rising after the push stops). What a programme builds therefore outlasts it, and the size of the gain depends on how long the money is sustained. These experiments use the stock channel of section 7.3, and under the flow channel nothing remains once the money stops.

Which pillar. The ranking is set by which domain is most persistent in a bloc, and not by spillovers. On the long fit the people pillar, now the more persistent of the two fast domains, accumulates most in every bloc at seven years, followed by the economy in the Mediterranean and the East and by place in the Continental and Nordic regions. At twenty-five years the order is the same in the Mediterranean and the East, while in the Continental and Nordic regions place overtakes people. The shorter fit put the economy first in the Mediterranean, because the crisis window made the Mediterranean economy

look more persistent than it is. Between 13 and 32 per cent of a push to people or the economy arrives, after twenty-five years, through the other two domains. The people-first ranking fits Rodríguez-Pose and Fratesi (2004), for whom education and human capital were the only Objective 1 spending with significant medium-term returns. It also fits Álvarez, Barbero, Orea and Rodríguez-Pose (2026), who find across 230 European regions that the less-developed ones gain most from spending on education. The direction of the loop adds to the ranking without reversing it. In the Mediterranean a people push reaches the economy at 21 per cent of its effect after twenty-five years, while an economy push reaches people at 5 per cent. In the struggling tier the economy does not feed people at all (section 6.2), so money that reaches only the economy cannot build people where they are scarcest.

Transfers. Through the imported money term, a seven-year transfer of one percentage point of GDP a year, about seven times the median intensity actually paid (section 3), raises the economy pillar of an eligible region off the absorption floor by 0.033 of its standard deviation by year seven. It raises that of a near-eligible region by 0.025 and that of a region on the floor by nothing. Under the people-stock measure fourteen eligible regions are on the floor, and under the upper-secondary measure nine are. The same machinery gives the number a policymaker would want. This is the constant intensity, in points of GDP a year above the European average and held for twenty-five years, that would lift each of the core seven above the 2008 bottom-quarter line by 2048. At the imported effect and gradient, no intensity up to fifty points a year lifts six of the seven, and Melilla would need 38.4. At average absorption the need runs from 5.7 points a year for Melilla to 25.5 for Sicily. For instance, Calabria would need 22.6 points on top of the 1.1 per cent it receives, some twenty times its actual intensity, every year for a generation, and at double the imported effect these figures halve. Sequencing a people push before the transfer on the floor regions adds at most 0.0004 of D to the two alone, both at the estimate on the bloc transitions and when the pooled transition is redrawn from the measurement bootstraps. No sequencing gain therefore exists in this system at the imported effect size.

Where to act. In a linear system with region intercepts, a change in conditions has the same long-run effect whether delivered by a regional or a national instrument. The weight of national against regional action is therefore the national share of the pull (section 6.1), which the experiments take as given. The data also show that 30 per cent of below-median regions in 2008, and 11 per cent in 2023, sit in countries where no region reaches the European median, so no national benchmark exists for them to converge to.

The experiments taken together. Three things survive every variant. The first is that relative positions move only when money is concentrated, because money spread at the average intensity produces no gain relative to Europe. The second is that gains accumulate with the years a programme is sustained and persist after it ends, so the movement the low-standing regions need takes decades of continuous money. The third is that absorption rises steeply with human capital. The regions cohesion policy exists for are the ones least able to use it, on either measure of capacity for the poorest of them. Cohesion policy as practised, spread widely, run in seven-year cycles and judged by its effect on output, would suit a system that adjusts quickly, absorbs money anywhere and moves in relative terms when everyone gets the same, and the system estimated here does none of those things.

9. Discussion

9.1 What the ecosystem is

The conceptual companion proposed three domains that support one another. The empirical companion found the configuration and its national scale but saw only one of the links from one year to the next. In this paper I find the system at the speed it operates, keeping its domains in proportion within a generation and changing its level over half a century or more. The earlier results are consistent with that reading. For instance, the weak annual links are the fast modes observed over too short a horizon. The euro crisis signature, the economy falling alone, is an imbalance in the Eastern system, where the economy has a fast mode of its own, and in large part a level shock in the Mediterranean system, where the slow direction loads most on the economy. For this reason the same hit corrected faster in the East than in the South. The decade-long recovery the empirical companion read as regions climbing back was mostly the national shocks reversing, which is what a slow mode would look like in the companion’s data. The national scale is also the slow mode’s origin, since the level a region is pulled toward is set by its country, and I would argue that the basin the companions describe is that national level. The region is where the imbalances live and where cohesion money is delivered, while the country is where the level is set.

9.2 What the claims rest on

Each of the paper’s claims rests on a test. The modes come from a transition whose estimator recovered known values on the long design, with 90 per cent intervals that covered the truth in 77.1 per cent of checks, just short of the 80 set in advance. The six pathways improve the log posterior by 82.9 over a model with the same intercepts, inputs and shocks and no links between pillars. A century of regional income data outside the model also shows less than half of a gap in level closing within a working life. The money effect, by contrast, did not recover and is imported. The model beats persistence out of sample in every year and the network from 2021, and it falls short of a country trend by 0.005, which misses the criterion set in advance. The crisis decomposition reproduces the observed gap as its implementation check and passes the pandemic test it was not tuned to. The low-standing core holds across every transition, condition and threshold tried, and an index built on a different logic agrees in direction, though more weakly. The money argument is arithmetic about relative intensity plus the discontinuity evidence, reported across a range that includes zero. The paper does not claim a causal pathway matrix, a steady state, or an estimate of what a euro does, and it says why not and what would supply them.

9.3 Why not a neural network, and why not a network of regions

I tried the first and considered the second. A neural network given the same inputs forecasts 2021 to 2023 worse than the ecosystem model, matches it in 2020 and does barely better than “nothing changes” (Figure 1), and neither matches a country trend. A flexible learner therefore finds no forecastable structure in these regions that the linear ecosystem misses, and it would give up the pathways, the decomposition and the projections. A network between regions, where neighbours influence one another, is absent from the model by design. Neighbours share a country and the country shock absorbs what a spatial link would explain, so the two are not separately identified on this panel. I ran the residual test that is possible, regressing the regional shocks the model leaves on the neighbours’ composite the year before (queen contiguity, 194 regions with a neighbour, leaving

out 16 islands and outermost regions without one, country-year means removed). This gives -0.0014 per standard deviation, with an interval of -0.0062 to +0.0017, against a shock standard deviation of 0.040. Beyond a region's own state, its country and the model, its neighbours' movements last year do not predict what happens to it this year. Figures 1 and 2 show the within-region network, the three domains and their links, and the between-region network is the natural extension.

9.4 An institutional reading of the blocs

The blocs were defined by geography and the model gives them no institutional content, although they coincide closely with groupings that comparative political economy has long described. The Continental and Nordic regions belong to the conservative and social-democratic welfare regimes of Esping-Andersen (1990) and to the coordinated market economies of Hall and Soskice (2001). The Mediterranean bloc is the southern model of Ferrera (1996), with an income-maintenance system that is generous to insiders and thin for everyone else, and with the family filling the gap. The Eastern bloc corresponds to the dependent market economies of Nölke and Vliegenthart (2009), whose growth has come through foreign direct investment (Bohle and Greskovits, 2012). On this reading the direction of the loop in each bloc has a plausible institutional counterpart. In the East, investment and employment arrive first and skills follow, which is the economy leading people. In the South, where the economy is slowest to correct an imbalance, I would argue that the slow adjustment reflects economies that had grown on domestic demand and lost the exchange rate as an instrument inside the monetary union (Johnston and Regan, 2016), so that an economic shock is worked off through a long compression of wages and employment. The missing feedback from the economy to people in the struggling tier fits a welfare state that invests little in the capabilities of those outside secure employment. I tested this reading by refitting the transition with countries grouped by regime in place of geography, on the same data and priors. Four regimes, social-democratic (Denmark, Finland and Sweden), conservative, southern and post-socialist, fit better than the three geographic blocs by 19.2 log-posterior units for nine further parameters, a gain per parameter comparable to the one the blocs bring over the pooled fit. The gain comes from separating the Nordic regions from the Continental ones, which geography had merged. In the social-democratic regime the economy feeds people more strongly than anywhere else in Europe, at 1.56 against 0.52 in the other direction. An imbalance between the two is corrected with a half-life of under four years, against eight in the conservative regime, nine in the post-socialist one and twenty-one in the southern one. The conservative regime has the loop close to balance, at 0.50 from people to the economy and 0.56 from the economy to people. The southern and post-socialist regimes reproduce the Mediterranean and Eastern blocs, with people leading in the South (0.77 against 0.20) and the economy leading in the East (0.75 against 0.56). An economy hit alone keeps 53 per cent of its loss after twenty-five years in the southern regime, 34 in the post-socialist, 33 in the social-democratic and 25 in the conservative. With eighteen regions the social-democratic slow mode is not pinned down, and I report only its fast modes. Splitting the post-socialist group into the dependent market economies of the Visegrád countries and the remaining periphery adds only 4.7 units for nine parameters. The data therefore do not separate the two, and the Visegrád regions show the Eastern pattern on their own (0.76 against 0.45). Among the three western regimes, the link from the economy to people is strongest in the social-democratic one and weakest in the southern one, and the speed at which an imbalance is corrected follows the same order. This is the order in which Esping-Andersen (1990) and Ferrera (1996) place them by the reach of social protection, so the loop the model estimates has an institutional counterpart that the data distinguish.

The regimes can also be replaced by a measured index. I let the transition vary linearly with a

national index of social investment, $B_c = B0 + z_c B1$, where z_c is fixed at its mean over 2000 to 2007 so that the crisis cannot move it. The index is the standardised average of three shares of GDP, namely social protection spending outside old age and survivors, education spending, and spending on active labour-market measures (Eurostat, 2026; European Commission, 2026b). With nine parameters it improves on the pooled fit by 30.0 log-posterior units, against 33.7 for the three blocs with eighteen parameters and 52.9 for the four regimes with twenty-seven. One standard deviation above the European mean, the link from the economy to people is 0.93 and an imbalance is corrected with a half-life of seven years. One standard deviation below it the link is 0.10 and the half-life twenty-one years, a difference in the link of 0.84 [0.35, 1.51], and the link from people to the economy moves the other way, from 0.49 to 0.88. An economy hit alone keeps 29 per cent of its loss after twenty-five years at the high value and 53 per cent at the low one. Education spending and working-age social protection carry the result, improving the fit by 27.3 and 26.3 units when entered alone, whereas active labour-market spending adds 8.5 and its interval includes zero.

The index is correlated with national income at 0.50, so I entered income per head in the same way. Income alone improves the fit by 6.3 units and moderates neither link, with a difference in the economy-to-people link of 0.20 [-0.45, 1.04]. With both entered together the fit improves by 43.2 units, and the difference associated with the index, at mean income, is 0.87 [0.38, 1.55]. Richer countries do not have a different loop for being richer, whereas countries that spend more on education and on the working-age population do. In my view this gives the regime result a measured counterpart, although the index is a national average over eight years and the association is not a causal estimate of what such spending does.

9.5 Limits

I see seven limits in the paper. Twenty-four years is short for a system whose slow mode is measured in decades. The fast modes are well determined, while the slow mode ranges from 44 to 389 years across twenty resamples, which is not a calibrated interval. The pathways are associations, and a design with single-pillar shocks followed over decades (university foundings, transport completions, eligibility changes) would test them as causes, although that is a different project. The effect of cohesion money is imported, so the paper's contribution on money lies in the arithmetic of relative intensity, the absorption floor and the map. The country shocks that carry most of the crisis are a residual. The evidence establishes only that the loss was national, not captured by spreads, by fiscal stance on either measure or by the housing and credit cycle, and not the regions' sectoral structure. The model forecasts relative positions slightly less well than a country trend. The pooled transition is an average of blocs whose loops run in different directions. Finally, projected low standing is a scenario whose core holds across everything tried, although its extent does not.

10. Conclusion

A European region is a slow ecosystem inside a national system. Its three domains keep each other in proportion within a generation, while its level is set by its country and has a half-life of fifty-five to seventy years. A century of regional income data outside the model agrees that less than half of a gap in level closes within a working life. The euro crisis hit the South's economy through its national layer. In a system whose level rides on the economy a hit of that kind still leaves 52 per cent of the loss after twenty-five years, against 32 per cent in the East, so that the part which came back did so mostly as the national shocks reversed. The pandemic left proportions alone, and the

South's small gain is consistent with the European money it received. Under the conditions Europe maintains, about a seventh of its regions stay at the bottom, seven of them whatever is assumed, and the independent trap index agrees in direction, though more weakly. Cohesion money could not have moved them, because spread at the average it buys no relative gain, on the causal evidence the poorest regions absorb it least, and the movement they need takes decades of sustained money. To lift Calabria above the line by 2048 would take, at average absorption, about twenty-three points of its GDP a year for twenty-five years, and under the floor that the imported gradient implies in its strict form, no sum at all. None of that is an argument against cohesion policy, and in my view it indicates what kind of policy would fit the system. Such a policy would be concentrated, sustained past a cycle, and aimed at capacity where the economy cannot yet take the money. The ecosystem the companions described can be estimated, and I would argue that a policy for Europe's regions has to be designed around its speed.

References

- Álvarez, I. C., Barbero, J., Orea, L. and Rodríguez-Pose, A. (2026). How institutions shape the economic returns to investment in European regions? *Economic Modelling*, 155, 107445.
- Barone, G., David, F. and de Blasio, G. (2016). Boulevard of broken dreams. The end of EU funding (1997: Abruzzi, Italy). *Regional Science and Urban Economics*, 60, 31 to 38.
- Bastolla, U., Fortuna, M. A., Pascual-García, A., Ferrera, A., Luque, B. and Bascompte, J. (2009). The architecture of mutualistic networks minimizes competition and increases biodiversity. *Nature*, 458(7241), 1018 to 1020.
- Beaumont, M. A., Cornuet, J.-M., Marin, J.-M. and Robert, C. P. (2009). Adaptive approximate Bayesian computation. *Biometrika*, 96(4), 983 to 990.
- Becker, S. O., Egger, P. H. and von Ehrlich, M. (2010). Going NUTS: the effect of EU Structural Funds on regional performance. *Journal of Public Economics*, 94(9 to 10), 578 to 590.
- Becker, S. O., Egger, P. H. and von Ehrlich, M. (2013). Absorptive capacity and the growth and investment effects of regional transfers: a regression discontinuity design with heterogeneous treatment effects. *American Economic Journal: Economic Policy*, 5(4), 29 to 77.
- Becker, S. O., Egger, P. H. and von Ehrlich, M. (2018). Effects of EU Regional Policy: 1989 to 2013. *Regional Science and Urban Economics*, 69, 143 to 152.
- Bohle, D. and Greskovits, B. (2012). *Capitalist Diversity on Europe's Periphery*. Ithaca: Cornell University Press.
- Bollen, K. A. (1989). *Structural Equations with Latent Variables*. New York: Wiley.
- Conway, A. (2026a, in preparation). A Regional Development Ecosystem Framework: Place, People and Economy in European Regions and Their National Systems.
- Conway, A. (2026b, in preparation). Two Crises, One Structure: Coupling and Scale in European Regional Ecosystems.
- Crescenzi, R. and Giua, M. (2020). One or many Cohesion Policies of the European Union? On the differential economic impacts of Cohesion Policy across member states. *Regional Studies*, 54(1), 10

to 20.

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.

Doz, C., Giannone, D. and Reichlin, L. (2011). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. *Journal of Econometrics*, 164(1), 188 to 205.

Eklund, G., Sibilía, 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.

Esping-Andersen, G. (1990). *The Three Worlds of Welfare Capitalism*. Cambridge: Polity Press.

European Commission (2026a). AMECO, the annual macro-economic database of the Directorate-General for Economic and Financial Affairs, chapter 17, cyclical adjustment of public finance variables (file dated 2 June 2026). Brussels: European Commission.

European Commission (2026b). Labour market policy database, expenditure by type of action (table `lmp_expsumm`). Brussels: Directorate-General for Employment, Social Affairs and Inclusion.

European Commission, Joint Research Centre (2021). ARDECO, the Annual Regional Database of the European Commission, 2021 vintage. Ispra: Joint Research Centre.

Eurostat (2026). Regional and national statistics, dissemination database: social protection expenditure by function (`spr_exp_func`), general government expenditure by function (`gov_10a_exp`), GDP per head (`nama_10_pc`) and the macroeconomic imbalance procedure indicators for private credit flow, private debt, the current account and deflated house prices (`tipspc20`, `tipspd20`, `tipsbp20`, `tipsho10`). Luxembourg: Eurostat.

European Union (2016). Consolidated version of the Treaty on the Functioning of the European Union. *Official Journal of the European Union*, C 202, 7 June 2016, 47 to 199.

Ferrera, M. (1996). The ‘Southern model’ of welfare in social Europe. *Journal of European Social Policy*, 6(1), 17 to 37.

Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S. and Tatlow, H. (2021). A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour*, 5(4), 529 to 538.

Hall, P. A. and Soskice, D. (eds) (2001). *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage*. Oxford: Oxford University Press.

Hampton, S. E., Holmes, E. E., Scheef, L. P., Scheuerell, M. D., Katz, S. L., Pendleton, D. E. and Ward, E. J. (2013). Quantifying effects of abiotic and biotic drivers on community dynamics with multivariate autoregressive (MAR) models. *Ecology*, 94(12), 2663 to 2669.

- 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.
- Ives, A. R., Dennis, B., Cottingham, K. L. and Carpenter, S. R. (2003). Estimating community stability and ecological interactions from time-series data. *Ecological Monographs*, 73(2), 301 to 330.
- Johnston, A. and Regan, A. (2016). European monetary integration and the incompatibility of national varieties of capitalism. *JCMS: Journal of Common Market Studies*, 54(2), 318 to 336.
- Martin, R. (2012). Regional economic resilience, hysteresis and recessionary shocks. *Journal of Economic Geography*, 12(1), 1 to 32.
- Martin, R. and Sunley, P. (2015). On the notion of regional economic resilience: conceptualization and explanation. *Journal of Economic Geography*, 15(1), 1 to 42.
- Nölke, A. and Vliegenthart, A. (2009). Enlarging the varieties of capitalism: the emergence of dependent market economies in East Central Europe. *World Politics*, 61(4), 670 to 702.
- Pellegrini, G., Terribile, F., Tarola, O., Muccigrosso, T. and Busillo, F. (2013). Measuring the effects of European Regional Policy on economic growth: a regression discontinuity approach. *Papers in Regional Science*, 92(1), 217 to 233.
- 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. (2026). Institutions at the helm of industrial policy. *SERIEs*. <https://doi.org/10.1007/s13209-026-00341-4>
- Rodríguez-Pose, A. and Fratesi, U. (2004). Between development and social policies: the impact of European Structural Funds in Objective 1 regions. *Regional Studies*, 38(1), 97 to 113.
- Rohr, R. P., Saavedra, S. and Bascompte, J. (2014). On the structural stability of mutualistic systems. *Science*, 345(6195), 1253497.
- Rosés, J. R. and Wolf, N. (eds.) (2019). *The Economic Development of Europe's Regions: A Quantitative History since 1900*. London: Routledge.
- Rosés, J. R. and Wolf, N. (2020). Rosés-Wolf database on regional GDP, version 6. Berlin: Humboldt-Universität zu Berlin.
- Saavedra, S., Reed-Tsochas, F. and Uzzi, B. (2009). A simple model of bipartite cooperation for ecological and organizational networks. *Nature*, 457(7228), 463 to 466.
- Toni, T., Welch, D., Strelkowa, N., Ipsen, A. and Stumpf, M. P. H. (2009). Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6(31), 187 to 202.

Appendix A. Reproducibility

The data build is in `paper3/v3/build_inputs.py` (measurement panel, money, national conditions) and `paper3/v3/pull_raw.py` (raw pulls with sidecars). The estimator is `paper3/v3/model.py`, with likelihood and smoother checks in `check_model.py` and `check_smoother.py`, recovery tests in `test_recovery.py`, `test_recovery_import.py` and `test_recovery_estimated.py`, fits

and validation in `run_phase_c.py` and `continue_fits.py`, and experiments in `phase_d.py` and `phase_e.py`. Revision checks are in `revision/*.py` (the transition by institutional regime `revision/d2_byregime.py`, by a measured institutional index `revision/d4_moderator.py`, further national conditions `revision/d5_shock_inputs.py`, the uncoupled benchmark `revision/d6_uncoupled.py` and `revision/d6b_fullse.py`, the refit on the cyclically adjusted fiscal impulse `revision/d1_capb_refit.py`, sectoral face of the country shock `revision/c5_sector_face.py`, neighbours `revision/c6_neighbours.py`, lifting intensities `revision/c7_intensity.py`, migration `revision/c8_migration.py`). The long panel is built by `paper3/v3/build_inputs_long.py`, fitted by `revision/l1_fit_long.py`, tested for recovery by `revision/l1_recovery_long.py` and fitted by bloc by `revision/l1_bybloc_long.py`, with results in `out_long/`. The environment switches `ECO_Y0` (start year), `ECO_CLEAN` (data folder), `ECO_REF` (reference set) and `ECO_OUT` (output folder) select the long design, and at their defaults the estimator reproduces the published fit. The baseline of this paper adds `ECO_FULLSE=1`, which lets regional shocks correlate across pillars, with the cyclically adjusted fiscal impulse in `clean_capb`, and its results are in `out_fse/`, produced by the chain `revision/f_run.sh`. The outside checks on the slow mode are in `code/analysis/slow_mode_outside.py` (output `output/slow_mode_outside.json`). The theoretical model is `paper3/eco/core.py`, with calibration in `abc_fsetargets.py`, region positions from `build_long_mapping.py` and experiments in `p_experiments.py` and `v21.py` (run with `ECO_LONG=1`, chained by `run_on_abcfse.py`). Every rule was written into the project ledger before its run, and the ledger records each withdrawal with its reason.

Appendix B. The nonlinear form and the feasibility domain

The ecosystem in generalised Lotka-Volterra form is $dX/dt = \text{kappa } S X (r_i + A X)$, with $S = \text{diag}(\text{eps}, 1, 1)$, A carrying -1 on the diagonal and the six pathways off it (a_{ij} the effect of stock j on the growth of stock i), and r_i a national vector plus a regional deviation. Its equilibrium is $X^* = -A^{-1} r$, its Jacobian $J = \text{kappa } S \text{diag}(X^*) A$, and its annual transition $\exp(J)$, and the fitted B of the body is the mean over regions of $\exp(J_i)$. I calibrated it by sequential Monte Carlo approximate Bayesian computation (Toni et al., 2009; Beaumont et al., 2009) to the fitted B of the long fit and to the pillar correlations and national share of the 203 regions observed in 2008, with tolerances from the data's own sampling error. The sampler reached its simulation cap before the planned tolerance, and the final population misses the cross-sectional targets by one and a half to two standard deviations and the larger links by one to two, always with correlations too high and links too weak, because strong mutualism in this form makes the stocks move together more than the 2008 cross-section shows. I therefore use the population only for the absorptive term below and for the national share of conditions cited in section 6.1. On it, the feasibility domain is 2.1 times (1.9 to 2.2) what it would be without the links. I measure the domain as the share of directions of the intrinsic-condition vector r , drawn uniformly on the sphere, for which all three stocks are positive at equilibrium, a share of one eighth without the links.

The absorptive term on the calibrated population. For each parameter set in the population, the 203 regions observed in 2008, 2013, 2019 and 2023 are placed at their 2008 positions as equilibria, and the eligible regions among them receive a seven-year transfer. A region with observed pillar scores z is placed at $X = \mu + s z$, where μ and s are the mean and standard deviation of each stock's equilibrium level in that parameter set's simulated cross-section, and simulated changes are converted back to observed standard deviations by dividing by s . These are 46 of the 51 eligible regions of section 7, because five Polish regions redrawn in the NUTS revisions have no 2008 position. Θ_0 and β are solved so that the mean change in the economy pillar matches Becker, Egger

and von Ehrlich (2013) for eligible regions with people scores above and below the European mean. These targets are 0.0465 and 0.0085 of the economy pillar's 2008 standard deviation, the 1.1 and 0.2 percentage points a year mapped through GDP's share of the composite. Both moments are reproduced for the whole population, with beta at 1.24 (1.10 to 1.98). On that population 39 per cent of eligible regions (35 to 48) are on the floor under the people-stock measure. When each receives the transfer alone, eligible and near-eligible regions in the East gain 0.023 of the economy pillar's standard deviation per seven-year transfer, those in the Continental bloc 0.029 and those in the Nordic bloc 0.107, while the median Mediterranean region gains nothing. Eligible and near-eligible regions absorb about the same, 0.018 in both cases, with the near-eligible ahead in three fifths of the population under the stock channel and in none of it under the flow channel. On the shorter fit's 152 regions the near-eligible absorbed twice as much, and the earlier population run on the present regions gives the present result, so the change comes from which regions are mapped rather than from the recalibration. The mean effect over the 46 is 0.018 against 0.068 for Becker et al.'s 2010 average, whose sample was mostly the EU15 before enlargement. Under the upper-secondary measure the fitted gradient is 0.51 (0.46 to 1.76), the mean effect 0.036, and the floor largely disappears, holding three of the 46 at the median gradient. With the below-average effect set to zero, 74 per cent of eligible regions are on the floor, although the moments can then be matched for only three fifths of the population, and with it at 0.5 points one region is. Under the stock channel 80 per cent (75 to 83) of a seven-year gain remains seven years after the money ends, and under the flow channel nothing does, by construction. These are the numbers section 7 refers to.