Zouhair Ait Benhamou
Université Le Havre Normandie, EDEHN, France
E-mail: zouhair.ait-benhamou@univ-lehavre.fr

LOCAL EFFECTS OF ENERGY DECOUPLING IN EU-27 ON LABOR: A SPATIAL INVESTIGATION

Original scientific paper
UDK: 338.2:331.5:502.131.1(4-672EU)
JEL classification: Q52, R11, R15, J21, C31
https://doi.org/10.17818/DIEM/2026/SI5
Paper received: June 24, 2026
Paper accepted: August 25, 2026

Abstract

The final push toward the EU's 2030 CO_2 target is likely to concentrate employment losses in energy-dependent regions. Using a spatial econometric decomposition on 1,131 NUTS-3 regions (2000–2022), we quantify regional employment exposure to this remaining decarbonization burden by separating two channels: the labor-output and output-CO_2 elasticities. Estimating a Spatial Durbin Model with a 2000 base-year Bartik shift-share instrument, we obtain an employment-CO_2 elasticity of $\varphi = 0.169$ ($SE = 0.0025$). Labor-output elasticities are statistically indistinguishable across regions, so the vulnerability gap is driven predominantly by sectoral carbon intensity. We report gross employment exposure – not net losses, since green-job creation is not modelled – under two scenarios: a uniform cut (15.5M jobs, 7.3%) and one differentiated by raw CO_2 (19.0M jobs, 8.9%). The estimates offer a quantitative basis for targeting Just Transition Mechanism support toward regions whose exposure exceeds 2% of employment.

Keywords: decarbonization, employment elasticity, NUTS-3, just transition

1. INTRODUCTION

The European Green Deal commits the EU to 55% emissions reductions by 2030 relative to 1990, a legally binding target with profound regional consequences. Though the EU has already achieved 31% reduction between 1990 and 2022 (World Bank, 2024), with a further 35% cut to close in just 8 years[1]. Closing this gap is not a uniform adjustment across space. Some regions are well-positioned through their decoupling policies, while others face a cliff, as they depend structurally on carbon-intensive production for which low-carbon alternatives do not yet exist at scale. These are not incremental adjustments, they require phasing out economic activities that currently sustain employment in some of Europe's most economically vulnerable regions.

While the EU acknowledges the regional unevenness of these cuts, it lacks the fine-grained employment impact assessment needed to guide support efficiently. Which regions face the deepest adjustment risk, and is that risk driven by labor-market rigidity (workers hard to retrain) or by sectoral carbon lock-in (production structures dependent on emissions)? We address these questions by decomposing the employment-CO_2 elasticity through the chain rule $\varphi = \beta \cdot \gamma$, where β is the employment-output elasticity and γ the output-CO_2 elasticity. Because regional labor markets are spatially integrated through trade, commuting, and factor mobility – interdependence that standard panel estimators cannot capture – we estimate a panel Spatial Durbin Model (SDM)

[1] 2030 CO_2 levels should represent 45% of 1990 levels, hence $(1 - 0.31) \times (1 - 0.35) \approx 0.45$

This work is licensed under a Creative Commons Attribution 4.0 International License.

with queen-contiguity weights across all 1,131 NUTS-3 regions, separating direct from indirect (spillover) effects. The two channels call for different responses: labor rigidity calls for worker retraining, while carbon lock-in calls for industrial investment and green-technology deployment.

The contribution is threefold: a chain-rule decomposition that attributes the employment-CO_2 elasticity to its labor-market (β) and production-structure (γ) sources, with delta-method standard errors; a Spatial Durbin Model on NUTS-3 panel data ($N = 25{,}990$) with formal Moran and Lagrange-multiplier specification search (Anselin, 1988; Anselin et al., 2026); and a Bartik shift-share instrument built from base-year 2000 sectoral employment shares interacted with post-2000 national sectoral GVA changes (Borusyak et al., 2022).

The structure of the paper is laid out as follows: the second section reviews strands of literature on the topic of labor elasticity to energy decoupling, its local effects and quantitative methods designed to capture spatial effects. The third section lays out the conceptual framework through a two-step estimation approach. The fourth section describes the data and NUTS vintage concordance procedure. The fifth section presents exhaustively the empirical strategy including the Bartik instrument construction, spatial model selection, LeSage & Pace (2009) decomposition, and reports the main results as well. Section six extends the analysis to the number of jobs potentially affected by the Fit-for-55 targets. Finally, section seven discusses policy implications and concludes.

2. LITERATURE REVIEW

The question of whether environmental regulation destroys jobs has generated a large and contested literature. Early quasi-experimental work centered on the United States: Vona et al. (2018) exploit variation in the stringency of the Clean Air Act across US metropolitan areas and find no detectable impact on aggregate employment, but a significant compositional shift toward engineering and technical skills. This "null aggregate, significant compositional" pattern became a reference point for the European debate. The EU Emissions Trading System (ETS), launched in 2005, offers the most studied European natural experiment. Dechezleprêtre et al. (2023) use installation-level inclusion thresholds in a difference-in-differences design to show that the ETS cut emissions by about 10% over 2005–2012 with no significant effect on firm profits or employment. Using comprehensive administrative microdata, Colmer et al. (2025) similarly estimate emission reductions of 14–16% with no significant impact on firm employment, output, or capital, as regulated firms responded through energy-saving investment rather than outsourcing or contraction. These firm-level nulls, however, mask regional and sectoral heterogeneity that becomes visible only in finer-grained spatial analysis.

Beyond the ETS, a second strand examines the broader employment consequences of the transition. Maldonado et al. (2024) map labor-market exposure across European regions using granular labor-force surveys and find that workers in high-carbon jobs lacking decarbonization options have fewer and less transversal skills, which limits their mobility into growing occupations – so the aggregate null may conceal large distributional effects across skill groups. Grashof & Basilico (2024) show that mono-sectoral (brown) regions face greater difficulty diversifying into green technologies than economies with a balanced makeup. Saucedo-Acosta & Carrillo-Ovando (2026) give this diversification argument a measurable form: using a fixed-effects panel of EU member states, they document a U-shaped relationship between the Economic Complexity Index and CO_2 emissions, robust to substituting fossil-fuel use or energy intensity. The emissions content of a productive structure is thus non-monotonic in its sophistication. This matters directly for our decomposition: if the output–emissions elasticity γ depends non-monotonically on productive structure, it cannot be treated as homogeneous across regions—precisely what motivates our spatially heterogeneous specification. Two qualifications limit direct transposition of their result: it is estimated at the member-state level, so it cannot speak to the within-country dispersion that drives regional exposure, and it is *aspatial*, allowing no role for the cross-regional spillovers that Vagnini et al. (2024) document for industrial emissions.

The regional dimension of these distributional effects is examined by Rodríguez-Pose & Bartalucci (2024), who show that many lagging EU regions are both more vulnerable to climate change and more exposed to the employment costs of the Green Deal, with some Southern European regions facing losses above three percent of the workforce linked to the exit from fossil-fuel activities. This concentration of adjustment costs in already-fragile regions is the political-economy rationale for the Just Transition Mechanism, though the authors rely on NUTS-2 data, which sits uneasily with the far more localized territories selected for JTF support. Benner et al. (2026) make a related point, arguing that just-transition outcomes are rooted in place-based structural, ideological, and economic factors. At the sector level, McDowall et al. (2023) map employment vulnerability across EU NUTS-2 regions using PRIMES decarbonization scenarios for four carbon-intensive industries (electricity, oil refining, steel, cement), finding vulnerability concentrated in a small number of regions that overlap the EU's coal and transition regions, though their indicator approach does not estimate causal elasticities. Related notions of labor-market vulnerability and exposure to decoupling are also documented by Borsekova & Korony (2023).

Sectoral employment concentration overlaps with the geography of stranded assets – assets that lose value prematurely under climate policy. Coal regions are the clearest illustration: Oei et al. (2020) document Germany's hard-coal phase-out from 1950 to 2018, over which employment fell from roughly 600,000 to fewer than 10,000 workers. Despite decades of transition support, regional unemployment remained persistently elevated and alternative job growth never fully offset mining losses, especially for older and lower-skilled workers – evidence that sectoral transition is not a temporary adjustment cost but deep structural change. Marin and Vona (2021), studying French manufacturing over 1997–2015, find that energy-price increases reduce employment in energy-intensive sectors while raising it in less carbon-intensive ones, so the net effect depends on regional sectoral composition. Since carbon pricing operates partly through energy prices, our output-emissions elasticity γ captures part of this channel at the regional level. On the methodological side, two-way fixed-effects (TWFE) panel estimation remains the workhorse for exploiting within-region variation while controlling for regional and year heterogeneity (e.g. Naqvi, 2021).

Claeys et al. (2024) distinguish aggregate employment effects from job reallocation, arguing that decarbonization can be employment-neutral in aggregate while inducing substantial sectoral and regional reallocation – which motivates region- and sector-specific analysis of exposure. For identification, the Bartik shift-share instrument, introduced by Bartik (1991) and formalized by Goldsmith-Pinkham et al. (2020), generates plausibly exogenous variation by combining initial local exposure (shares) with aggregate shocks external to local units (shifts); it is numerically equivalent to using base-period industry shares as instruments. Le Gallo & Kamarianakis (2011) use a related technique to separate industry mix from structural regional differences in NUTS-2 productivity, and Borusyak et al. (2022, 2025) show that valid inference requires either exogenous shifts or exogenous shares, with practical guidance on each case. Spatial econometrics enters because regional labor markets are not isolated: Furková & Chocholatá (2021) estimate Spatial Durbin Models for employment across 259 EU NUTS-2 regions and find significant indirect effects of neighbors' GDP and compensation, and their later work (Furková & Chocholatá, 2023) confirms positive spatial autocorrelation of regional employment rates.

The closest spatial-econometric antecedent to ours is Vagnini et al. (2024), who apply spatial panel methods to NUTS-2 industrial CO_2 over 2008–2020 and find significant positive spatial autocorrelation in emissions, with negative direct and spillover effects of R&D and education. We extend their framework in three directions: working at NUTS-3 rather than NUTS-2, focusing on employment rather than emissions, and using a chain-rule decomposition that separates the labor-output from the output-emissions channel. Giannakis & Mamuneas (2022) use a Cobb-Douglas production function with SDM estimation to study regional labor-market resilience in EU NUTS-2 regions and find strong spatial spillovers in the productivity response to shocks. Our decomposition extends this by introducing emissions as an intermediate term in the chain, attributing the employment-CO_2 elasticity to its structural sources rather than treating it as a reduced-form black

box. This structural view implies that spatial heterogeneity in emissions intensity should translate into differentiated labor-market exposure: the elasticity is unlikely to be uniform across territories, but shaped by persistent regional characteristics, among which geographic location is central.

Finally, Demeterova (2024) argues that most regional-development strategies view just transition through a narrow lens of economic performance (competitiveness, income, growth distribution) and calls for a shift toward learning-oriented goals – a reminder that the material conditions analyzed above do not by themselves define a successful transition.

3. CONCEPTUAL FRAMEWORK

This paper grounds its estimated employment elasticity to energy decoupling in a structural production function approach, that formalizes the link between environmental constraints and labor market outcomes through a chain-rule decomposition of the labor-emissions elasticity. We assume the existence of a regional production function characterized by constant returns to scale Cobb-Douglas setup, where output denotes Y and is a function of capital K, labor L, and greenhouse gas emissions CO_2, the latter serving as a proxy for carbon-intensive energy inputs. Let β denote the labor-output elasticity, and γ output-emissions elasticity. By applying the chain rule such that $\varphi = \beta \cdot \gamma$, we isolate whether a region's transition risk is driven by labor market dependencies or the carbon intensity of its industrial base.

Parameter φ is defined by the non-linear transformation $\varphi = f(\beta, \gamma)$ which we define as a vector of our estimates. The variance of this product is approximated by the quadratic form $G'\Sigma G$ where Σ is the variance-covariance matrix of the estimators and G is the gradient vector of the transformation with respect to each parameter β and γ. Specifically, the gradient vector G is composed of the partial derivatives, namely $\frac{\partial \varphi}{\partial \beta} = \gamma$ and $\frac{\partial \varphi}{\partial \gamma} = \beta$. Under the common empirical assumption that the two elasticities are estimated from independent models or that their error terms are asymptotically uncorrelated, the covariance term $\mathrm{cov}(\beta, \gamma)$ is retained because both coefficients are co-determined by output Y. In this case, the variance of the chain-rule elasticity is expressed as:

$$V(\varphi) = \left(\frac{\partial \varphi}{\partial \beta}\right)^2 V(\beta) + \left(\frac{\partial \varphi}{\partial \gamma}\right)^2 V(\gamma) + 2\left(\frac{\partial \varphi}{\partial \beta}\right)\left(\frac{\partial \varphi}{\partial \gamma}\right)\mathrm{cov}(\beta, \gamma)$$

Substituting the partial derivatives into the equation yields the practical formula:

$$V(\varphi) = \gamma^2 \sigma_\beta^2 + \beta^2 \sigma_\gamma^2 + 2\gamma\beta\mathrm{cov}(\beta, \gamma)$$

The resulting standard error, $SE(\varphi) = \sqrt{V(\varphi)}$ is critical for the paper's inferential validity. It accounts for the propagation of uncertainty through the chain: a region might have a precisely estimated labor-output relationship (β), but if the link between emissions and output (γ) is measured with significant noise, perhaps due to the downscaling of EDGAR data or the volatility of GVA in specific NUTS-3 regions, the resulting employment risk φ will carry a wider confidence interval.

4. DATA AND SUMMARY STATISTICS

We combine macroeconomic data for 1,131 NUTS-3 regions in 27 countries over the 2000–2022 period, drawn from the EU's ARDECO database. Four main variables are used to that effect: Gross Value Added (GVA) in real terms, Annual hours worked, Local workforce and Capital stock in real terms. In order to account for demographic differences, relevant variables are expressed on a per worker basis. Note that CO_2 emissions are not provided at NUTS-3 level by the EDGAR database. To that effect, we make the assumption that CO_2 emissions are proportional to sectoral GVA, with sector-specific intensities. Thanks to the rich NACE classification provided by ARDECO, we can overcome potential biases arising from too coarse a sectoral classification. All macroeconomic variables are expressed on a per-worker basis to account for demographic and labor force

differences across regions. Specifically, per-capita normalization ensures that regional comparisons reflect productivity and intensity metrics rather than aggregate size effects. CO_2 emissions data are available at the NUTS-2 regional level only. To derive NUTS-3-level estimates, we employ a sectoral downscaling approach that takes into account the heterogeneity in sectoral composition across regions. The method rests on the assumption that within each NUTS-2 region, CO_2 emissions are distributed across sectors in proportion to that sector's economic activity, with sector-specific carbon intensities, namely:

$$CO_2(\text{NUTS-3}) = \sum_{j=1}^{10} CO_2(\text{Sector}_j) \frac{\text{GVA Sector}_j}{\sum \text{GVA NUTS-3}}$$

This approach ensures that: (a) NUTS-3 estimates aggregate back to the original NUTS-2 totals by construction, and (b) regional variations in CO_2 reflects both sectoral composition differences and sectoral decarbonization trends at the NUTS-2 level. Beyond employment headcount, we incorporate two additional labor market indicators from ARDECO to assess job quality: real compensation per hour worked and real compensation per employee. These allow us to distinguish between wage-driven and hours-driven changes in total labor income. We also track nominal unit labor cost (compensation per unit of output), which captures the wage-productivity balance and signals regional competitiveness dynamics. These metrics are essential for evaluating whether regions achieving decoupling simultaneously create quality employment or experience jobless growth. The dataset spans the full 2000–2022 period, allowing us to analyze dynamics across distinct policy regimes: the pre-financial crisis period (2000–2008), the recovery and early climate policy phase (2008–2015), and the post-Paris Agreement implementation of EU climate targets (2015–2022). We also stratify analyses by region type (coastal vs. inland) and country to capture institutional and sectoral heterogeneity. Coastal NUTS-3 regions are identified via standard GISCO geographical definitions. Our sample consists of 1,131 NUTS-3 regions observed over 23 years. The median region has approximately 115,000 workers and €62,400 GVA per capita (2022). Regional heterogeneity is substantial: the top-decile region by GVA per worker is approximately 3.2 times richer than the bottom-decile, reflecting the well-documented East-West productivity gradient in Europe. These disparities are central to our analysis, as we investigate whether convergence in GVA per worker over 2000–2022 has occurred alongside, or at odds with convergence in carbon intensity and employment creation. The set of maps below reports the spatial distribution for three of the main indicators brought to bear in our analysis, namely local GDP (Gross Value Added, GVA), hours worked and CO_2 emissions. All variables are expressed per worker in order to take into account demographic differences.

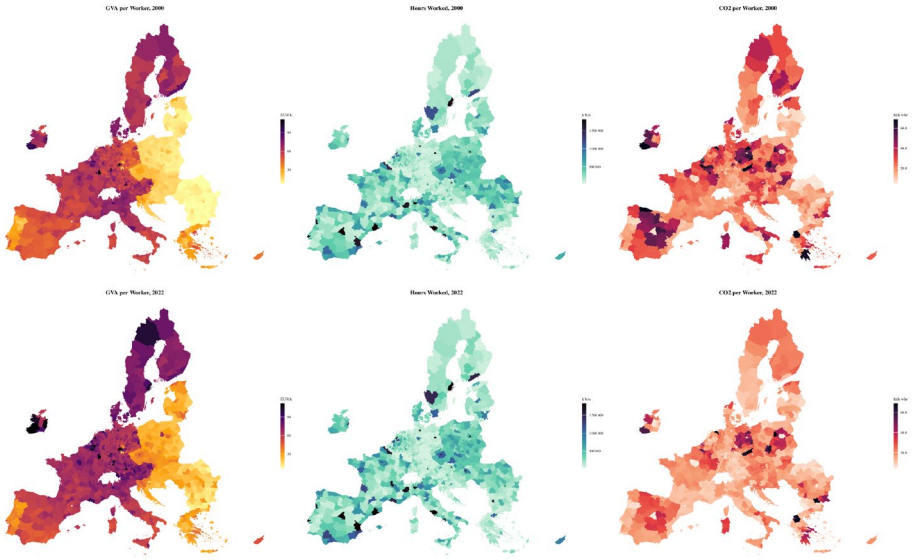


Figure 1 Spatial distribution of GVA per worker, hours worked per worker and CO_2 emissions per worker across NUTS-3 regions, 2000 and 2022

The first row reveals a striking pattern of regional convergence in labor productivity. In 2000, Central and Eastern European (CEE) regions display substantially lower value added per worker relative to Western European cores (UK, France, Benelux, Germany). By 2022, this gap has compressed noticeably. CEE regions have exhibited genuine productivity gains, even if an absolute productivity hierarchy persists. By contrast, the hours-worked maps tell a story of remarkable structural persistence. Major metropolitan clusters remain focal points of employment concentration in both 2000 and 2022. Closer inspection suggests a secondary shift, though: the observed concentrations appear slightly less extreme in 2022, with marginally more diffusion toward secondary cities and surrounding regions.

The most dramatic transformation appears in carbon intensity per worker. The 2000 map shows stark geographic clustering around coal-dependent regions of Poland, Benelux, France and Germany's industrial heartlands, and parts of Southern Europe. These contrast sharply with lighter Scandinavian and Western zones. By 2022, the clustering has diffused substantially across the continent. The spatial unevenness suggests that carbon reduction may have been a function of energy policy, technological investment, and sectoral composition, which are advantages concentrated in wealthy, early-adopting regions.

The maps in figure 1 therefore document a two-speed Europe when it comes to reliance on fossil energy and its use in output: one achieved simultaneous productivity growth, employment diffusion, and decarbonization (mostly the West and North), and another in transition, with productivity catching up but employment remaining concentrated and carbon intensity lingering above Western peers. The decoupling between CO_2 emissions and economic activities that EU- and nationwide policies have called for has not been achieved uniformly.

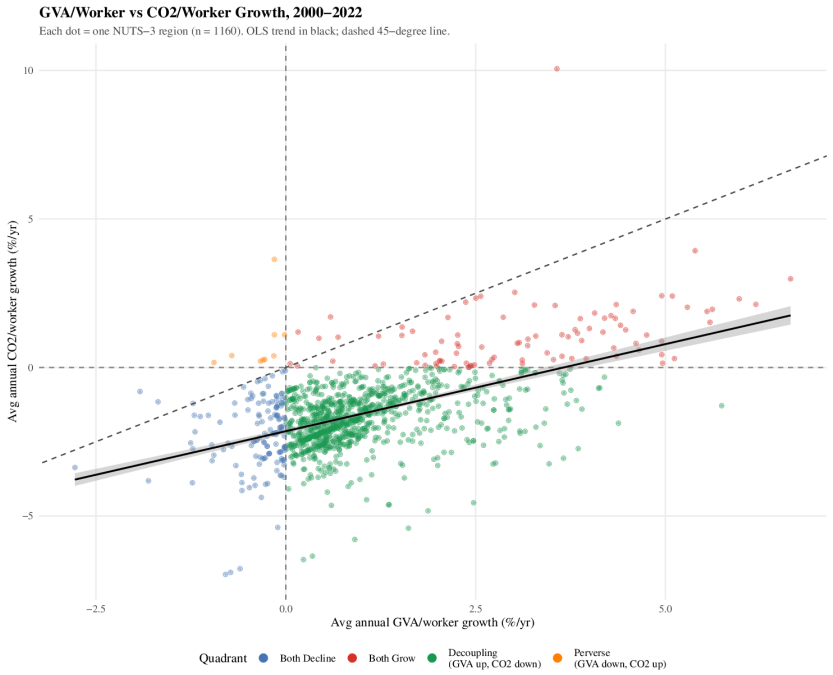


Figure 2 Bivariate classification of NUTS-3 regions by GVA and CO_2 emissions growth per worker, 2000–2022

This bivariate classification reveals a geographically stratified sustainability transition. The continent is sharply divided between regions that have achieved true decarbonization alongside growth and those still locked in carbon-intensive expansion, or worse, contraction without emissions reduction. Most regions have achieved a certain level of decoupling, as green areas suggest. Red regions form a concerning aggregate across mostly scattered industrial zones. These areas have achieved the productivity gains documented in the first map's convergence narrative, but without

shedding their fossil fuel dependence. This represents the classic *development trap*: growth in manufacturing-heavy or energy-intensive sectors that rely on coal or gas, or economies where sectoral composition has not shifted toward services. Blue spots appear sporadically, often in post-industrial zones and declining rust-belt peripheries. Counterintuitively, these represent a form of unintended decoupling which illustrate the concept of *managed decline* (Eaton, 2021) where regions that have deindustrialized, shed both jobs and emissions simultaneously. Though economic decline is not a sustainability strategy, it does achieve the arithmetic of lower emissions. Finally, orange/yellow zones are rare but concerning where they appear. These represent the worst outcome: regions losing economic output while their per-capita emissions rise, suggesting either a shift toward dirtier economic structures, inefficient industrial transition, or incomplete data. These should trigger policy scrutiny. Overall, the figure shows that for those outliers, there is a heterogenous level of spatial autocorrelation, which suggests that decoupling – or lack thereof – does not happen in a vacuum, and that further investigation is needed. Decoupling in that sense is an achieved outcome, concentrated in regions with high institutional capacity, deep capital reserves for energy transition, and sectoral advantages in services and advanced manufacturing. The persistence of red areas suggests that EU convergence and climate neutrality operate on fundamentally different spatial logics, and may require targeted policy intervention beyond market mechanisms alone. As a result, the scatterplot in figure 2 illustrates the need to distinguish between absolute and relative decoupling. Top right quadrant describes a muddled state, where both emissions and economic activities growth, though it can be argued that most locations record higher economic growth relative to CO_2 emissions, thus achieving relative decoupling. Bottom right quadrant depicts absolute decoupling, i.e. CO_2 emissions decline as economic growth is sustained. It is by far the most dominant regime across NUTS-3 locations. Left-side quadrants report those locations that have failed to achieve decoupling, either because of managed decline, or because energy intensive activities have located themselves there with little to no effect on economic activities *per se*.

The correlation matrix in figure 3 shows that decoupling remains incomplete even in dynamic terms. Indeed, the most striking result is the persistent positive correlation between GVA and CO_2 growth, even though all variables are expressed as growth rates rather than levels. The conversion to growth rates removes spurious correlation from common stochastic trends and isolates genuine dynamic relationships. This means that regions that grow their per-capita GVA fast also, on average, grow their per-capita CO_2. It reflects a structural feature of regional economic trajectories: growth and emissions expansion remain mechanically linked, independent of the underlying level of capital or employment. In addition, a strongly negative correlation between capital and employment growth appears across all regions. This encodes a critical dynamic, namely, regions that grow capital stock rapidly tend to shed employment. This output-employment decomposition reveals that Europe's growth story since 2000 is fundamentally one of capital substitution for labor, most acute inland and in productivity-led regions, with a temporal inflection in 2008 (when firms responded to crisis through labor shedding) and subsequent institutional entrenchment of part-time work. The bright spot, that coastal and employment-led regions exist, is shadowed by the evidence that these predominantly lack decoupling, creates an apparent policy *trilemma*: inclusive growth (jobs), green growth (decoupling), and productive growth (efficiency) have proven difficult to achieve simultaneously across Europe's regions.

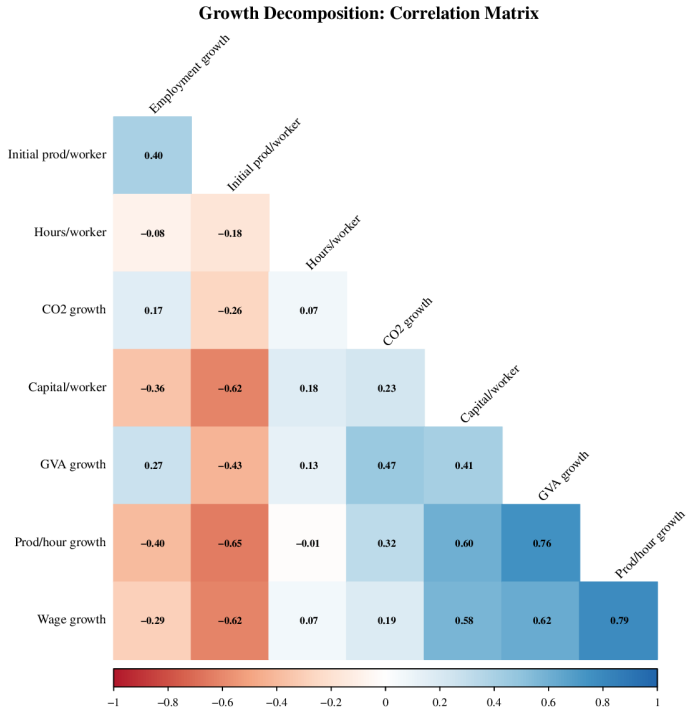


Figure 3 Correlation matrix in growth levels

Decoupling analysis alone provides an incomplete picture of regional sustainability and inclusive growth. While documenting the relationship between output and emissions is essential for climate policy, it tells us little about whether regions achieving decoupling simultaneously create or destroy employment opportunities. This matters critically for territorial cohesion: a region may successfully reduce its carbon footprint, but if this is achieved through capital-deepening and labor displacement rather than employment expansion and productivity gains shared with workers, the climate victory becomes a social failure. Conversely, regions experiencing rapid employment growth may be widening economic opportunities for workers while locking in fossil-fuel dependency, which is a cohesion gain shadowed by environmental liability.

5. SPECIFICATIONS AND RESULTS

Identifying the causal effect of CO_2 emissions reductions on regional employment presents a conceptual and methodological puzzle. Emissions and employment are jointly determined outcomes reflecting regional economic structure, sectoral composition, historical development patterns, and policy regimes. A simple correlation between these variables would conflate multiple mechanisms operating simultaneously, like labor market adjustment, capital stock evolution, technological adoption, and sectoral reallocation, making causal interpretation impossible. To untangle these mechanisms and provide meaningful estimates for policy analysis, we decompose the focal parameter of interest into two estimable components via the chain rule. This decomposition not only improves identification but also illuminates distinct economic channels through which decarbonization affects employment, enabling policymakers to target interventions toward the appropriate lever (labor market support versus sectoral transition assistance).

We assembled a balanced panel of 1,131 NUTS-3 regions (EU-27 excluding 21 isolated island regions without contiguous neighbors) observed annually from 2000 to 2022, yielding 24,882 region-year observations. Using Eurostat GISCO boundaries, a queen contiguity spatial weight matrix is row-standardized and yields an average 5.2 neighbors per region. The Bartik shift-share

instrument for addressing endogeneity in γ estimation was constructed from base-year (2000) sectoral employment shares interacted with national sectoral GVA growth rates.

φ the employment-emissions elasticity is the focal parameter quantifying regional employment exposure to decarbonization. This decomposition mirrors the productivity accounting tradition established by Solow (1957) and extended by Jorgenson and Griliches (1967), but applies it to environmental constraints rather than growth accounting. Instead of decomposing output growth into its labor and capital contributions, we decompose emissions reduction impacts into labor market adjustment channels (β) and sectoral production constraints (γ). The advantage of this decomposition is twofold: methodologically, it allows distinct causal identification strategies appropriate to each parameter (instrumental variables for γ where simultaneity is problematic, panel fixed effects for β where within-region variation suffices), and substantively, it provides transparency about which mechanism dominates employment losses: are jobs lost because firms cannot produce without emissions, or because workers cannot transition to available alternatives?

Selecting the empirically appropriate specification requires making multiple methodological choices simultaneously: choice of estimator (OLS, fixed effects, random effects, instrumental variables), specification of spatial structure (aspatial, spatial lag, spatial error, or spatial Durbin), and choice of time dimension (cross-sectional averaging or panel). Rather than imposing a single specification *ex ante* based on prior assumptions, we employ a sequential model selection procedure guided by formal statistical tests and information criteria. This approach begins with a candidate set of specifications spanning different estimator families and spatial structures, then eliminates candidates that fail diagnostic tests or are nested within and dominated by superior alternatives. The result is a selection process that is transparent and justified at each elimination stage. The model selection funnel proceeds through three successive stages. The first addresses the fundamental question of unobserved heterogeneity in the labor-output relationship: does each region have its own persistent employment-output relationship (arguing for fixed effects), or is a pooled specification adequate? Moreover, does the output-CO_2 relationship suffer from simultaneity bias requiring instrumental variables, or does OLS suffice? Stage 2 addresses spatial structure: do neighboring regions' outcomes affect own-region outcomes through demand and factor linkages (spatial lag), are residuals spatially correlated due to omitted spatial variables (spatial error), or both? Stage 3 compares candidate spatial models using log-likelihood and information criteria, selecting the specification that best balances fit and parsimony while satisfying diagnostic adequacy tests. This sequential logic embodies the principle that simpler, more parsimonious models should only be abandoned when statistical evidence compels otherwise, which guards against overfitting and produces specifications that are robust across data subsets and time periods. We begin by estimating the labor-output elasticity β using standard linear panel regression models. The fundamental specification regresses employment on output:

$$\ln L_{it} = \alpha + \beta \ln Y_{it} + \upsilon_{it}$$

Where subscript i indexes across NUTS-3 regions in the EU, t indexes years (2000 to 2022), and υ_{it} denotes the error term. The parameter of interest β captures the elasticity of employment with respect to output, the central quantity for understanding labor market responsiveness to regional growth. The naïve Pooled OLS assumes a constant intercept α across all regions and years, treating Europe as a single integrated labor market with uniform employment-output relationships. This specification however ignores profound differences in regional labor market structures. These persistent regional characteristics, with geographic advantages, historical sectoral specialization, institutional features of labor relations, educational infrastructure, are time-invariant yet fundamentally determine how a given output change translates into employment change. Ignoring these differences by pooling all regions and years would leave them in the error term, potentially biasing estimates if they are correlated with output growth. The Pooled OLS specification (M1) is compared against four other specifications: Two-Way Fixed Effects (TWFE) and spatial specifications, namely SEM, SAR and SDM in panel form. As a result the Hausman test decisively rejects the pooled OLS specification ($p < 0.001$), confirming that region-specific and time-specific fixed effects are

essential. The Hausman specification test compares the pooled OLS estimator to the fixed effects (FE) estimator under the null hypothesis that region-specific unobserved heterogeneity is uncorrelated with output growth. The Hausman test statistic is:

$$H = \left(\hat{\beta}_{\mathrm{FE}} - \hat{\beta}_{\mathrm{OLS}}\right)' \left[\mathrm{Var}\left(\hat{\beta}_{\mathrm{FE}} - \hat{\beta}_{\mathrm{OLS}}\right)\right]^{-1} \left(\hat{\beta}_{\mathrm{FE}} - \hat{\beta}_{\mathrm{OLS}}\right)$$

Which is distributed as χ^2 under the null hypothesis. For our data, the Hausman test yields a value of $H = 847$ ($p < 0.001$). The within-region identification strategy TWFE isolates the effect of annual output fluctuations on employment changes, controlling for permanent regional characteristics and common macroeconomic shocks. β elasticity is estimated at $\hat{\beta} = 0.298$ which indicates that a 10% increase in regional output is associated with a 2.98% increase in regional employment. Modern economies rely on productivity improvements (more output per worker) rather than proportional labor-input growth. However, TWFE assumes spatial independence: a region's employment-output relationship operates in isolation from neighbors' outcomes. This assumption is violated when regions are economically integrated through supply chains, commuting, and factor market linkages. To that effect, we apply Anselin's Lagrange Multiplier tests to TWFE residuals to diagnose the source and nature of spatial dependence: LM-lag and LM-error, and both reject handily the null hypotheses, confirming that spatial dependence operates through both channels: (1) neighboring regions' employment levels directly affect own-region employment (spatial lag), and (2) unobserved regional characteristics that affect own-region employment are correlated across space (spatial error). This dual mechanism suggests the Spatial Durbin Model (SDM), which nests both structures, is appropriate.

Finally, maximum likelihood estimation was applied to all five model specifications (M1 to M5) on the same sample of 24 ,882 observations. Information criteria balance model fit against complexity for both Akaike and the Bayesian criterion, with results reported below:

Table 1 AIC/BIC selection model, β estimate

Model	AIC	Δ-AIC	BIC	Δ-BIC
M5: SDM-Panel	−89,894	Baseline	−89,853	Baseline
M4: SAR-Panel	−86,662	3,232	−86,638	3,215
M3: SEM-Panel	101,024	190,918	101,048	190,901
M2: TWFE	−76,509	13,385	−76,484	13,369
M1: Pooled OLS	30,972	120,866	31,004	120,857

The SDM-Panel (M5) is overwhelmingly preferred by both AIC and BIC, with Δ-AIC of 13,385 relative to TWFE (M2). This massive margin indicates that ignoring the Durbin WX term (spatially lagged output, θ) and the spatial lag (ρ) introduces severe misspecification that far outweighs any penalty for the additional parameters. The AIC/BIC dominance of SDM is the decisive criterion for model selection. While the Anselin tests confirmed both lag and error autocorrelation exist, the information criteria make clear that the Durbin specification provides the superior fit tradeoff.

Table 2 Estimated values for β per model

Model	Coef β	SE	Student t-stat
M1	1.064	0.009	119.528
M2	0.298	0.020	14.871
M3	0.326	0.004	78.649
M4	0.189	0.003	54.198
M5	0.344	0.004	79.643

We then proceed with the output/CO_2 elasticity, which is fundamentally endogenous: regions experiencing positive technology shocks simultaneously reduce emissions intensity and increase output, creating downward bias in OLS estimates. To address this, we employ the Bartik shift-share instrument developed by Borusyak et al. (2022). Regional base-year shares are equilibrium objects, so identification leans on the shocks (EU-wide ETS/directive/technology forces) being external, and the leave-one-out removes own-country feedback. The foundational insight is that EU climate policy operates primarily at the sectoral level, not the regional level. The Emissions Trading System sets economy-wide carbon prices and sectoral caps; renewable energy directives mandate percentage renewable targets; energy efficiency standards apply to product classes. A region specializing in coal mining inevitably faces steeper policy-driven emissions reduction pressure than a region specializing in services, simply by sectoral exposure. Moreover, this differential exposure is plausibly exogenous to the region, since it depends on EU-level policy decisions and the region's pre-existing sectoral specialization, not on any region's current productivity shocks or growth dynamics.

$$\text{Bartik}_{it} = \sum_{s} \text{Share}_{is,2000} \times \Delta\text{lnGVA}_{s,t}^{\text{National}}$$

Here, $\text{Share}_{is,2000}$ is region i's employment share in sector s in the base year 2000 (a predetermined variable, established before any future regional shocks), and $\text{GVA}_{s,t}^{\text{National}}$ is the national-level change in sectoral GVA in year t.

The instrument is thus a weighted average of national sectoral trends, where weights are the region's historical specialization. A region with 30% of its 2000 workforce in coal mining and 50% in services would have an instrument heavily influenced by coal sector national emission trends (which declined steeply post-2008 as coal plants closed across Europe) and less influenced by service sector trends (which remained flatter). Critically, the base-year specialization is established in 2000, before any future shocks, ensuring that the instrument is predetermined and exogenous to future regional output shocks. The logic behind the Bartik instrument is that regional output growth in 2010–2015 for instance cannot causally affect the region's 2000 sectoral composition (past determines present, not vice versa). Therefore, variation in regional exposure to post-2000 national sectoral decarbonization, as captured by the Bartik instrument, is exogenous to any future regional productivity or growth shocks. The instrument provides leverage on regional CO_2 emissions that is independent of regional output, precisely the source of variation needed to identify γ causally.

We proceed with the first stage by regressing logged CO_2 emissions on the Bartik instrument, namely:

$$\text{lnCO2}_{it} = \alpha_i + \tau_t + \pi \cdot \text{Bartik}_{it} + u_{it}$$

Coefficient $\hat{\pi}$ measures the conditional association between regional emissions and exposure to national sectoral decarbonization, after partialling out the controls and fixed effects. Its precision is summarized by the cluster-robust (Kleibergen–Paap) first-stage F, which tests instrument *relevance* only: it asks whether the instrument's partial correlation with lnCO_2 is strong enough for the 2SLS estimator to have acceptable finite-sample properties, and carries no information about the exclusion restriction. Because the design is just-identified, no test of instrument validity is available; exogeneity is maintained as an identifying assumption and defended on the substantive grounds set out above. We benchmark relevance against the Montiel Olea–Pflueger (2013) effective F rather than the conventional rule of thumb, and report tF-adjusted confidence intervals (Lee et al., 2022) for the second-stage elasticity, since the Stock–Yogo critical values are derived under homoskedasticity and are not valid with clustered standard errors.

Table 3 First-stage diagnostics for the Bartik instrument (endogenous variable $lnCO_2$; GVA-downscaled NUTS-3 emissions)

	Fixed-effects structure	
2-3 Diagnostic	Country×year	Region + year
	(primary)	(textbook)
First-stage coefficient $\hat{\pi}$	-4.146	0.948
Cluster-robust SE (NUTS-3)	(0.292)	(0.088)
t-statistic	-14.2	10.7
Cluster-robust first-stage F (Kleibergen–Paap)	202.1	115.1
Cragg–Donald F	2,120.7	1,540.8
Partial R^2 (instrument)	0.074	0.055
Weak-IV threshold (Stock–Yogo, 10%)	> 10	
Wu–Hausman endogeneity (p)	0.003	—

Notes: First stage $\ln\mathrm{CO}_{2,it} = \pi\,\mathrm{Bartik}_{it} + x'_{it}\delta + \mathrm{FE} + u_{it}$, where CO_2 is NUTS-2 EDGAR sectoral CO_2 allocated to NUTS-3 by sectoral GVA (conserving NUTS-2 totals). Controls $x = \ln K, \ln Y$, NACE-A10 sector shares, coastal dummy. Standard errors clustered on NUTS-3. The single-instrument, just-identified design admits no over-identification test. The instrument is strong under both fixed-effects structures ($F = 202$ and $F = 115$, far above the Stock–Yogo threshold), and Wu–Hausman rejects exogeneity of $\ln\mathrm{CO}_2$ ($p = 0.003$) in the primary specification.

Table 3 reports the diagnostics. The instrument is a strong predictor of downscaled emissions and enters with the expected sign: in the primary country×year specification $\hat{\pi} = -4.15$ (cluster-robust $\mathrm{SE} = 0.29$, $t = -14.2$), so regions historically specialized in sectors on steeper EU decarbonization trajectories record larger relative emission declines. The cluster-robust (Kleibergen–Paap) first-stage F is 202 and the Cragg–Donald F exceeds $2{,}000$, both far above the Stock–Yogo weak-instrument threshold of 10; the instrument explains roughly 7% of the conditional variation in $\ln\mathrm{CO}_2$ (partial $R^2 = 0.074$). Relevance is robust to the fixed-effects structure: under textbook region+year effects the first stage remains strong ($F = 115$, partial $R^2 = 0.055$, $t = 10.7$). With a single instrument the model is just-identified, so no over-identification test is available; the Wu–Hausman test rejects exogeneity of $\ln\mathrm{CO}_2$ ($p = 0.003$), confirming that instrumentation is warranted. We note that part of this strength reflects the construction of the emissions measure: because the downscaled series is a smooth function of sectoral GVA and NUTS-2 emission structure – the same primitives that enter the shift-share – the instrument tracks it more mechanically than it would a directly measured NUTS-3 series. We therefore read the first-stage F as establishing statistical relevance rather than as independent corroboration of the identifying variation, and rest identification on the shift-exogeneity argument above.

Because the shift-share instrument is by construction one-dimensional – a single index aggregating predetermined sectoral exposure against common national shocks – the model is exactly identified, and no over-identification test applies to it directly. One can artificially over-identify the model by partitioning the instrument into sector-group sub-instruments, but the resulting Sargan–Hansen test does not assess the identifying assumption. Under the shift-share framework of Borusyak et al. (2022), identification rests on the exogeneity of the national sectoral shocks, and the estimand is a shock-weighted average of heterogeneous output–emissions elasticities. When those elasticities differ across sectors – the maintained premise of our heterogeneity analysis – sub-instruments loading on different sectors estimate different weighted averages, so an over-identification test rejects even under fully valid instruments (Angrist and Imbens, 1995; Parente and Santos Silva, 2012). We accordingly do not rely on over-identification testing, and defend the exclusion restriction on the economic grounds that post-2000 national sectoral decarbonization was driven by EU-level policy external to any single region's labour-market dynamics.

Further specifications as for γ are offered: pooled OLS, TWFE and spatial models SEM, SAR and SDM in panel form. While the Anselin battery test suggests that there is spatial autocorrelation, it points to SEM as the better model candidate instead of SDM. Nevertheless, we stick with the latter due to its better AIC/BIC fit, as the table below reports:

Table 4 AIC/BIC selection model, γ estimate

Model	AIC	Δ-AIC	BIC	Δ-BIC
M5: SDM-Panel	−82,653	Baseline	−82,611	Baseline
M4: SAR-Panel	−72,120	10,533	−72,096	10,515
M2: TWFE	−63,226	19,427	−63,201	19,410
M3: SEM-Panel	110,204	192,857	110,229	192,840
M1: Pooled OLS	22,800	105,453	22,832	105,443

The SDM-Panel achieves the lowest AIC by a large margin compared to TWFE. This indicates that despite the Anselin tests suggesting error autocorrelation dominates, the SDM's flexibility in modelling both spatial lag and spatial error structures provides a substantially better fit.

Table 5 Estimated values for γ per model

Model	Coef	SE	Student t-stat
M1: Pooled OLS	0.288	0.004	80.708
M2: TWFE	0.372	0.021	17.906
M3: SEM-Panel	0.487	0.004	132.354
M4: SAR-Panel	0.147	0.003	50.199
M5: SDM-Panel	0.492	0.004	132.627

This illustrates an important methodological principle: model selection via AIC is preferred over hypothesis testing in the presence of large samples and complex specifications, it explicitly penalizes over-parameterization while rewarding fit improvements. An estimated output-CO_2 elasticity of $\hat{\gamma} = 0.492$ indicates substantial carbon lock-in but also meaningful scope for decoupling. A 10% reduction in CO_2 emissions is associated with a 4.92% reduction in regional output, all else equal. The spatial lag ($\rho = 0.654$) confirms that output in neighbouring regions is tightly linked, reflecting integrated supply chains and value-added networks across regions. Having estimated β and γ using identical spatial panel specifications, we combine them via the chain rule and compute its standard error using the delta method, which yields:

$$\hat{\varphi} = \hat{\beta} \times \hat{\gamma} = 0.344 \times 0.492 \;\rightarrow\; \hat{\varphi} = 0.169$$

and:

$$\sigma_{\varphi} = \sqrt{\gamma^2\sigma_{\beta}^2 + \beta^2\sigma_{\gamma}^2 + 2\gamma\beta\mathrm{cov}(\beta,\gamma)} \;\rightarrow\; \sigma_{\varphi} = 0.0025$$

We note that the covariance component is very close to zero, which is eminently positive for two reasons: first, it means that the standard error for the CO_2-employment elasticity is not understated if the covariance is eliminated through the assumption of independent values of β and γ. Second, it also means that the downscaling strategy of CO_2 emissions from NUTS-2 to NUTS-3 does not introduce a bias due to the common use of regional GVA.

The employment-CO_2 elasticity of 0.169 is our focal parameter. A 10% reduction in regional CO_2 emissions is associated with a 1.69% decline in regional employment, holding all other factors constant. In the context of the EU Fit-for-55 policy (55% CO_2 reduction by 2030), this translates to an aggregate employment impact of approximately $55\% \times 0.169 \approx 9.3\%$ on average, though with

substantial regional heterogeneity. In addition, this elasticity encompasses both direct labor-output responses and spatial spillovers built into the SDM structure through the ρ and θ coefficients.

Furthermore, the spatial setup allows us to use LeSage-Pace decomposition of direct and indirect effects embedded in the SDM. Unlike non-spatial models where the regression coefficient directly equals the policy effect, the SDM's spatial structure, through the lag term ρWy and the Durbin term θWX creates a more complex relationship between the independent variable and outcomes. A change in one region's CO_2 emissions affects not only that region's employment directly, but also neighboring regions' employment through spatial feedback loops (factor reallocation, trade adjustment, commuting patterns). Mathematically, these are derived from the reduced-form specification

$$y = (I_n - \rho W)^{-1}(X\Lambda + \theta WX + \varepsilon)$$

where the inverse of the spatial structure matrix $(I_n - \rho W)^{-1}$ captures how shocks propagate across regions. The direct effect averages the diagonal elements, the indirect effect averages the off-diagonal elements, and the total effect sums both. This decomposition allows us to assess whether the aggregate employment impact is driven primarily by own-region labor-output dynamics or by spatial spillovers from neighboring regions.

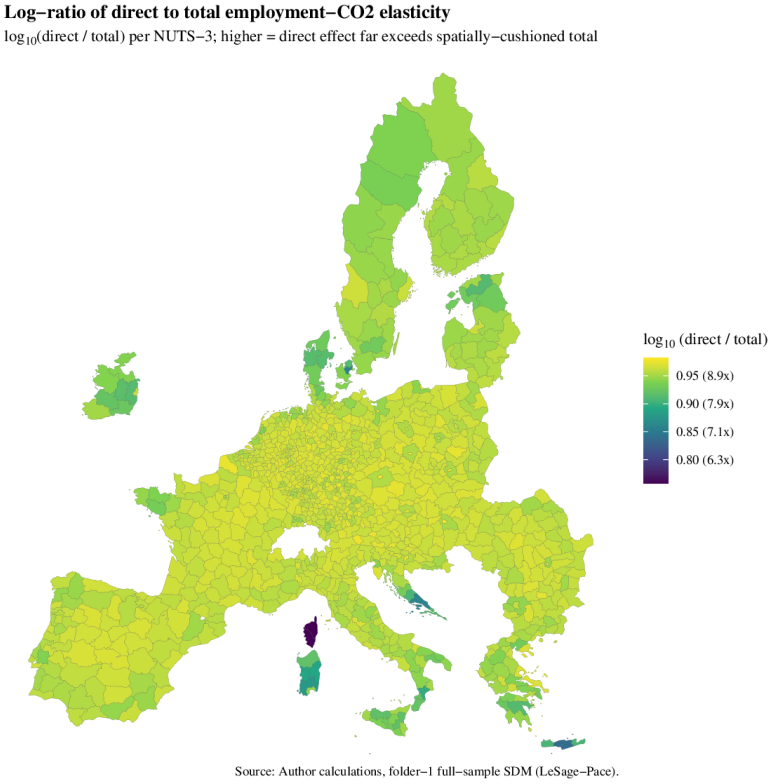


Figure 4 Spillover ratio of indirect to direct effects across EU NUTS-3 regions (LeSage-Pace decomposition)

The spillover ratio is negative and remarkably uniform across the EU. It is striking given the heterogeneity in regional size, economic structure, and network connectivity. The narrow range of 7.1 pp suggests a deeply embedded, EU-wide spatial equilibrium rather than region-specific idiosyncrasies. When a region cuts CO_2 emissions and experiences direct employment losses, the indirect (negative) spillover represents a gain in that region's employment through spatial reallocation. An average ratio of 89 means that for every 100 jobs lost directly to a region's own CO_2 cuts, approximately 89 jobs are recaptured through spillovers from neighboring regions. The underlying mechanisms for this recapture deal with trade reallocation, where cuts in emissions shift

demand from the original region to another one that supplies the same energy-intensive output. Similarly, labor is relatively mobile across regions, especially within national labor markets and Schengen area. Workers displaced from Region A's emission-intensive sectors migrate to neighboring regions with growing opportunities in non-emissions-intensive sectors. Spatial integration is so deep that unilateral emissions cuts in one region trigger substantial automatic reallocation. The negative sign (losses become gains through spillovers) reflects the zero-sum nature of this competition: one region's industrial job loss is another's gain at the regional scale.

Nevertheless, if all regions were to cut emissions (at varying degrees) then only 8.6% of total jobs are cushioned as there is no reallocation and/or mobility to be had. When all regions simultaneously cut CO_2 under Fit-for-55, the spatial reallocation benefits that individual regions receive are offset by the simultaneous cuts in their neighbors. Region A captures demand from its neighbors' reduced production, but Region A also cuts, denying its neighbors the same reallocation benefit. The net effect at the EU aggregate is near-zero. Furthermore, the 89% regional spillover ratio is uniform, but the direct elasticity φ varies by 17-fold across regions, -0.037 for Finland to 0.654 for Slovenia. This means that for every 100 jobs high-elasticity regions lose through the direct effect, neighboring regions gain roughly 89 jobs through spillovers, which implies moderate net losses, though they come at the price of large initial losses.

6. FIT-FOR-55 AND CO_2 EMISSION CUTS

The Fit-for-55 policy sets a binding target of 55% net greenhouse gas (GHG) emission reductions by 2030 relative to 1990 levels. Our empirical measure is energy-related CO_2, which accounts for the large majority of EU-27 GHG emissions; we treat the two as broadly comparable for the purpose of scaling employment exposure, while acknowledging that non-CO_2 gases follow distinct abatement paths and that the mapping is therefore approximate. The EU has already achieved substantial progress: emissions in EU-27 fell by 31% between 1990 and 2022. The remaining effort is thus a 24-percentage-point cut relative to the 1990 baseline, which, applied to the lower 2022 stock, corresponds to a reduction of approximately 35% from 2022 levels: if 2022 emissions stand at $(1 - 0.31) = 69\%$ of 1990 and the 2030 target is $(1 - 0.55) = 45\%$ of 1990, then the required cut from 2022 is $1 - 0.45/0.69 \approx 0.35$. Using our estimated employment-CO_2 elasticity $\varphi = 0.169$, the aggregate employment-weighted impact of this marginal effort is roughly 7.3%. This is substantially lower than the -13.5% aggregate impact implied by the full 55% cut from 1990 baseline, reflecting that much of the decarbonization burden has already been absorbed by the EU economy during 1990–2022.

We build two scenarios to evaluate the local effects of decarbonization on jobs. The first scenario assumes all NUTS-3 regions achieve the same proportional reduction from 2022 emissions. To ensure EU-27 achieves exactly 24 pp additional reduction, we simulate a uniform cut of 35% for all regions. This provides a baseline comparison assuming equal responsibility across all regions. A uniform cut is not a realistic assumption as it imposes a disproportionate burden on areas with low CO_2 emissions, and does not enact substantial cuts on those that emit the most. To close the remaining 24 pp emissions gap by 2030, a uniform 35% additional reduction from 2022 levels would expose approximately 15.5 million jobs (7.3% of current employment) on average. Regional variation reflects only country-level elasticity differences, not burden differentiation. The range of 26 pp (-24% to $+1\%$) is exactly proportional to the 55pp range seen in the full 55% cut scenario, but at 47% only of the magnitude.

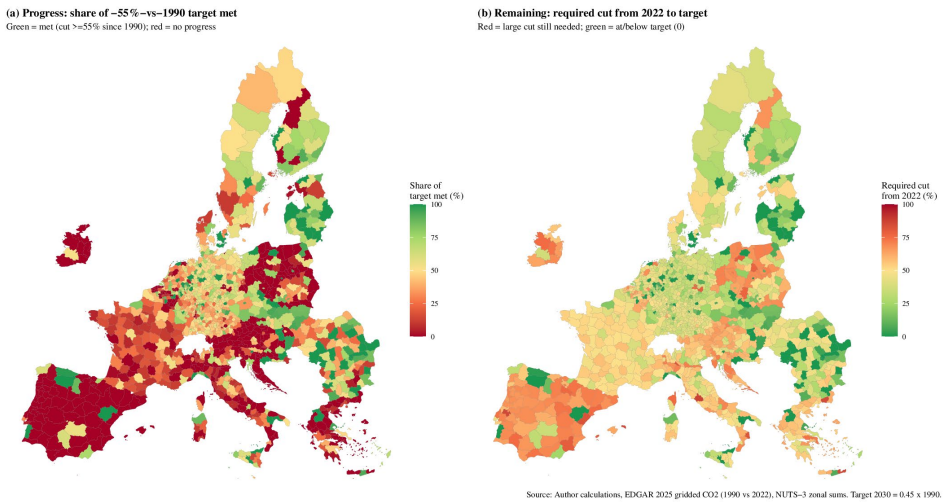


Figure 5 EU-27 remaining emissions reduction gap to the 2030 Fit-for-55 target

Table 6 Uniform scenario — 35% CO_2 emissions cut

Metric	Value	Interpretation
Total jobs at risk (thousands)	−15,503	7.3% of baseline employment
Mean impact	−7.27%	Average region exposure 7.3% of employment
Std. deviation	3.45%	Low variation due to identical cuts
Min loss	−24.4%	High-elasticity country region
Max gain	+1.37%	Modest gains due to negative elasticity

The second scenario offers an alternative differentiation, where regions cut proportionally to their absolute emissions (kt CO_2) from 2022. The formula writes:

$$cut_i = 0.35 + \mu \times \frac{\mathrm{CO}_{2_i} - \mathrm{CO}_2(\mathrm{National}_i)}{\mathrm{CO}_2(\mathrm{National}_i)}$$

where μ calibrated within each country in order to ensure the total sum of weighted-average emission cuts is equal to 35%. This approach targets large absolute emitters (metropolitan and industrial regions mainly) while potentially sparing small high-intensity regions. While a uniform cut affects some regions more than others, it also affects disproportionally countries over others. Economic development does not seem to factor into which countries is affected most in the aggregate, but rather the sectoral composition and its reliance on fossil energy.

Table 7 Differentiated cuts scenario – 35% CO_2 emissions cut

Metric	Value	Interpretation
Total jobs at risk (thousands)	−19,010	8.9% of baseline employment
Mean impact	−4.98%	Lower on average than CO_2/worker
Std. deviation	7.05%	Moderate variation
Min loss	−48.6%	Less extreme than per-worker
Max gain	+3.38%	Comparable gains as per-worker

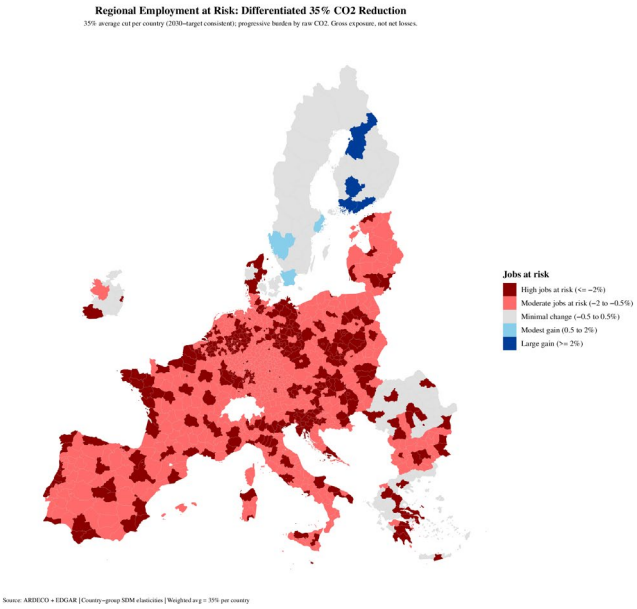


Figure 6 Regional employment impact classification under the differentiated cuts scenario

Table 8 Regional impact classification

Category	Employment Impact	Economic Interpretation
High exposure	$< -2\%$	Severe impact requiring intervention
Moderate exposure	-2% to -0.5%	Significant but manageable
Minimal Change	$\pm 0.5\%$	Near-negligible impact
Modest Gain	$+0.5\%$ to $+2\%$	Minor positive reallocation
Large Gain	$> +2\%$	Substantial employment creation

Differentiating by absolute emissions (in kt CO_2) is more lenient toward small high-intensity regions because their absolute emissions are low. However, it concentrates cuts on big industrial cities and areas, where employment is concentrated. This results in higher total jobs at risk despite lower mean loss, because the burden falls on high-employment metropolitan areas. The standard deviation of 7.05% is intermediate between uniform (3.45%) and per-worker (9.25%). The differentiated cuts scenario allows us to identify those areas most likely to be affected, and classify them by order of employment impact magnitude. Under this scenario, 82.8% of regions face employment losses. Unlike the full 55% scenario where 95% faced losses, the marginal 35% cut from 2022 allows more regions to experience near-negligible impacts. Only 19% in "Large Loss" (vs. 36.2% per-worker), while over half of regions experience "Moderate Loss," indicating a more balanced distribution. Fewer regions face catastrophic impacts, though aggregate jobs are at higher risk due to concentration in high-employment cities.

7. CONCLUSIONS

This paper quantifies the regional employment cost of the European Union's remaining decarbonization effort by decomposing the employment–CO_2 elasticity into a labor–output and an output–CO_2 channel, estimated on 1,131 NUTS-3 regions over 2000–2022 within a Spatial Durbin framework with Bartik shift-share instrumentation. The central estimate, $\varphi = 0.169$ ($SE = 0.0025$), implies that the 24-percentage-point cut still required between 2022 and 2030 carries an average employment exposure of roughly 7.3%, which is materially smaller than the -13.5% that a naïve

55%-from-1990 reading would suggest, because the bulk of the adjustment has already been absorbed by the European economy. The decomposition delivers the paper's main message. Vulnerability to decarbonization is not a question of labor-market flexibility, which is statistically indistinguishable across regions, but of carbon lock-in in regional output. This is why inland industrial regions face a burden 49% larger than coastal ones, and why the cost concentrates in fossil-dependent economies – Slovenia, Slovakia, Poland – irrespective of their level of development. The spatial structure adds a second-order qualification: while a region acting alone recaptures close to 89% of its direct job exposure through reallocation from neighbors, this channel collapses to roughly 9% of cushioning once all regions cut simultaneously, since every region is denying its neighbors the same relief. Unilateral and coordinated decarbonization are therefore not scalable versions of one another, and policy designed on the former will overstate the resilience available under the latter. These results provide a quantitative foundation for targeting Just Transition support. Rather than treating decarbonization losses as diffuse, our estimates identify the specific regions that exhibit high-φ and are fossil-dependent territories clearing the 2% employment-loss threshold, where the €100-billion-plus Just Transition Mechanism should concentrate. The choice between a CO_2-per-worker and a raw-CO_2 basis is shown to be a genuine distributional lever rather than a technical detail: the former protects metropolitan workforces at the cost of small industrial regions, the latter does the reverse. Making that trade-off explicit is, in our view, a precondition for a transition that is both credible and politically durable. Several limitations bound these conclusions. The analysis measures gross employment exposure and does not net out the green jobs created by Just Transition investment, so the estimates are best read as the scale of adjustment to be managed rather than a forecast of net losses. Elasticities are estimated at the country-group level for robustness and abstract from within-region sectoral and skill composition; the scenarios are log-linear approximations with cut rates clamped to a plausible range; and the 2022 emissions baseline is dictated by data availability rather than the 1990 reference of the policy target itself. Extending the framework to model green-investment offsets and to trace the sectoral reallocation of displaced workers is the natural next step, and the one most directly relevant to designing the compensation instruments this paper helps to target.

Author Contributions: The author confirms sole responsibility for the following: study conception and design, data collection, analysis and interpretation of results, and manuscript preparation.

Funding: The research presented in the manuscript did not receive any external funding.

Conflict of Interest: None.

Acknowledgements: The author thanks the participants of the Dubrovnik International Economic Meeting and the two anonymous referees for constructive comments that improved the manuscript.

Acknowledgement of AI or AI-assisted Tools Use: Artificial-intelligence tools (Claude, Anthropic) were used in a limited, supporting capacity only: for data cleaning and internal-consistency sanity checks; for assistance with coding corrections, dry runs to verify the coherence of results, and interpretation of error messages; and for light language editing (sentence structure, vocabulary, and clarity). All research questions, modelling choices, empirical strategy, results, and conclusions are the author's own, and the author takes full responsibility for the content of the paper.

REFERENCES

Angrist, J. D., & Imbens, G. W. (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American Statistical Association*, *90*(430), 431–442. https://doi.org/10.1080/01621459.1995.10476535

Anselin, L. (1988). Spatial econometrics: Methods and models. Dordrecht: Kluwer Academic Publishers. https://doi.org/10.1007/978-94-015-7799-1

Anselin, L., Serenini, R., & Amaral, P. (2026). Alternative perspectives on spatial econometrics model specification search. *Spatial Economic Analysis*, *21*(2), 288–335. https://doi.org/10.1080/17421772.2025.2521095

Bartik, T. J. (1991). Who benefits from state and local economic development policies? Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. https://doi.org/10.17848/9780585223940

Benner, M., Grünenwald, V., & Kastrup, J. (2026). The geography of just transitions: A place-based framework. Regional Studies, *60*(1). https://doi.org/10.1080/00343404.2025.2606364

Borsekova, K., & Korony, S. (2023). Resilience and vulnerability of regional labour markets: Principal component analysis of labour market efficiency in the EU. *Regional Studies*, *57*(12), 2373–2390. https://doi.org/10.1080/00343404.2022.2042507

Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, *89*(1), 181–213. https://doi.org/10.1093/restud/rdab030

Borusyak, K., Hull, P., & Jaravel, X. (2025). A practical guide to shift-share instruments. *Journal of Economic Perspectives*, *39(*1), 181–204. https://doi.org/10.1257/jep.20231370

Claeys, G., Le Mouel, M., Tagliapietra, S., Wolff, G. B., & Zachmann, G. (2024). The Macroeconomics of Decarbonisation: Implications and Policies. Cambridge University Press. https://doi.org/10.1017/9781009438353

Colmer, J., Martin, R., Muûls, M., & Wagner, U. J. (2025). Does pricing carbon mitigate climate change? Firm-level evidence from the European Union Emissions Trading System. *The Review of Economic Studies*, *92*(3), 1625–1660. https://doi.org/10.1093/restud/rdae055

Dechezleprêtre, A., Nachtigall, D., & Venmans, F. (2023). The joint impact of the European Union emissions trading system on carbon emissions and economic performance. *Journal of Environmental Economics and Management,* 118, 102758. https://doi.org/10.1016/j.jeem.2022.102758

Demeterova, B. (2024). Assessing regional capabilities-oriented approaches for European just and sustainable transitions. *Regional Studies*, *58*(8), 1530–1542. https://doi.org/10.1080/00343404.2023.2242894

Eaton, E. (2021). Approaches to energy transitions: Carbon pricing, managed decline, and/or green new deal? *Geography Compass*, *15*(2). https://doi.org/10.1111/gec3.12554

Furková, A., & Chocholatá, M. (2021). Spatial econometric approach to the EU regional employment process. *Central European Journal of Operations Research*, *29*(3), 1037–1056. https://doi.org/10.1007/s10100-020-00714-5

Furková, A., & Chocholatá, M. (2023). Spatial heterogeneity and spillovers of employment in the EU regions. Zagreb *International Review of Economics and Business*, *26*(2), 163–182. https://doi.org/10.2478/zireb-2023-0019

Giannakis, E., & Mamuneas, T. P. (2022). Labour productivity and regional labour markets resilience in Europe. *The Annals of Regional Science*, *68*(3), 691–712. https://doi.org/10.1007/s00168-021-01100-y

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review, 110*(8), 2586–2624. https://doi.org/10.1257/aer.20181047

Grashof, N., & Basilico, S. (2024). Divergent paths for green transition: Regional disparities in technology diversification. *Regional Studies*, *59*(1). https://doi.org/10.1080/00343404.2024.2363327

Jorgenson, D. W., & Griliches, Z. (1967). The explanation of productivity change. *The Review of Economic Studies*, *34*(3), 249–283. https://doi.org/10.2307/2296675

Le Gallo, J., & Kamarianakis, Y. (2011). The Evolution of Regional Productivity Disparities in the European Union from 1975 to 2002: A Combination of Shift–Share and Spatial Econometrics. *Regional Studies*, *45*(1), 123–139. https://doi.org/10.1080/00343400903234662

Lee, D. S., McCrary, J., Moreira, M. J., & Porter, J. (2022). Valid *t*-ratio inference for IV. *American Economic Review*, *112*(10), 3260–3290. https://doi.org/10.1257/aer.20211063

LeSage, J., & Pace, R. K. (2009). Introduction to spatial econometrics. Boca Raton, FL: Chapman & Hall/CRC. https://doi.org/10.1201/9781420064254

Maldonado, J. E., Vandeplas, A., Vanyolos, I., Vigani, M., et al. (2024). Assessing green job dynamics in the EU: A comparison of alternative methodologies. European Commission, Directorate-General for Economic and Financial Affairs. Luxembourg: Publications Office of the European Union.

Marin, G., & Vona, F. (2021). The impact of energy prices on socioeconomic and environmental performance: Evidence from French manufacturing establishments, 1997–2015. *European Economic Review*, *135*, 103739. https://doi.org/10.1016/j.euroecorev.2021.103739

McDowall, W., Reinauer, T., Fragkos, P., et al. (2023). Mapping regional vulnerability in Europe's energy transition: development and application of an indicator to assess declining employment in four carbon-intensive industries. Climatic Change, 176, 7. https://doi.org/10.1007/s10584-022-03478-w

Montiel Olea, J. L., & Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics, 31*(3), 358–369. https://doi.org/10.1080/00401706.2013.806694

Naqvi, A. (2021). Decoupling trends of emissions across EU regions and the role of environmental policies. Journal of Cleaner Production, 323, 129130. https://doi.org/10.1016/j.jclepro.2021.129130

Oei, P.-Y., Brauers, H., & Herpich, P. (2020). Lessons from Germany's hard coal mining phase-out: Policies and transition from 1950 to 2018. *Climate Policy, 20*(8), 963–979. https://doi.org/10.1080/14693062.2019.1688636

Parente, P. M. D. C., & Santos Silva, J. M. C. (2012). A cautionary note on tests of overidentifying restrictions. *Economics Letters*, *115*(2), 314–317. https://doi.org/10.1016/j.econlet.2011.12.047

Rodríguez-Pose, A., & Bartalucci, F. (2024). The green transition and its potential territorial discontents. *Cambridge Journal of Regions, Economy and Society*, *17*(2), 339–358. https://doi.org/10.1093/cjres/rsad039

Saucedo-Acosta, E. J., & Carrillo-Ovando, R. W. (2026). Linking Economic Complexity & the environment: the case of the European Union. *DIEM Dubrovnik International Economic Meeting*, *11*(1), 107. https://doi.org/10.17818/DIEM/2026/1.10

Solow, R. M. (1957). Technical change and the aggregate production function. *The Review of Economics and Statistics, 39*(3), 312–320. https://doi.org/10.2307/1926047

Vagnini, C., Canal Vieira, L., Longo, M., & Mura, M. (2024). Regional drivers of industrial decarbonisation: A spatial econometric analysis of 238 EU regions between 2008 and 2020. *Regional Studies, 59*(1). https://doi.org/10.1080/00343404.2024.2380369

Vona, F., Marin, G., Consoli, D., & Popp, D. (2018). Environmental regulation and green skills: An empirical exploration. *Journal of the Association of Environmental and Resource Economists, 5*(4), 713–753. https://doi.org/10.1086/698859

World Bank (2024). World Development Indicators: Carbon dioxide (CO2) emissions excluding LULUCF per capita. Washington, DC: World Bank.