

Revenue Portfolio Design, Fiscal Network Position, and Inequality: Evidence from Europe

Achillefs Karakotsios

akarakots@econ.auth.gr

Aristotle University of Thessaloniki

Thomas Panagiotou

Aristotle University of Thessaloniki

Constantinos Katrakilidis

Aristotle University of Thessaloniki

Research Article

Keywords: revenue portfolio, fiscal architecture, tax structure, income inequality, fiscal network, Shannon entropy, eigenvector centrality, dynamic network Durbin model, defactored instrumental variables, VAT compliance, Europe

Posted Date: October 6th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11077881/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Revenue Portfolio Design, Fiscal Network Position, and Inequality: Evidence from Europe

Achillefs Karakotsios*, Thomas Panagiotou, Constantinos Katrakilidis

School of Economics, Aristotle University of Thessaloniki, Greece

akarakots@econ.auth.gr · thompana@econ.auth.gr · katrak@econ.auth.gr

*Corresponding author: Achillefs Karakotsios (akarakots@econ.auth.gr)

Abstract

In the European Single Market, the distributional consequences of how a country taxes depend not only on its own revenue choices but on where those choices sit within the European fiscal network: the set of economies whose revenue structures most resemble the domestic one. Using OECD Revenue Statistics for 29 European economies over 1995 to 2022, we represent each tax system as a revenue-share vector over 28 positions and build a tax-similarity network from the cosine similarity of those vectors. Two indicators follow: tax diversification, the normalized Shannon entropy of a country's row in the network, and fiscal-network centrality, its eigenvector centrality. A dynamic network Durbin model estimated by defactored instrumental variables separates direct inequality diffusion from structural diffusion through fiscal-neighbour characteristics. Tax diversification is consistently associated with lower disposable-income inequality, fiscal-network centrality with higher inequality, and neighbours' VAT compliance gaps with higher inequality in the full sample, though the last association weakens after 2005. The two domestic results form an asymmetry: portfolio diversification is a lever national governments can pull alone, whereas network centrality is a structural exposure that only coordinated European action can relieve.

Keywords: revenue portfolio; fiscal architecture; tax structure; income inequality; fiscal network; Shannon entropy; eigenvector centrality; dynamic network Durbin model; defactored instrumental variables; VAT compliance; Europe

JEL Classification: H20, D63, C33, C36, F42

1. Introduction

Between 1995 and 2022, Finland and Romania both operated inside the European Single Market and faced the same body of EU fiscal directives. Yet their distributional outcomes diverged sharply: Finland remained among Europe's most equal economies and Romania among its least, a persistent disposable-income inequality gap of several Gini points. Standard explanations, such as welfare-state scale, labour-market institutions, and development levels, account for much, but not all, of this divergence. European economies spend considerable effort designing their tax systems, choosing how much to rely on personal income taxes, social security contributions, consumption levies, or property taxes. The standard view of this problem is essentially domestic: each country solves a national trade-off between revenue adequacy, efficiency, and redistribution, and the distributional outcome follows from that choice. This paper argues that the domestic view is incomplete. In the European Single Market, national tax systems are not chosen in isolation. They are embedded in a fiscal neighbourhood of structurally similar systems, and the distributional consequences of domestic tax design depend partly on who the neighbours are. Finland's tax system sits near the centre of the European fiscal-similarity network, which exposes it to competitive pressures that constrain redistributive manoeuvre. Romania's more idiosyncratic structure leaves it more room to deviate.

The empirical evidence supporting this argument is striking. Inequality is highly persistent across European economies, as would be expected given the institutional inertia of tax-benefit systems and wage structures. But what drives cross-country differences in the level and trajectory of inequality is less obvious. The aggregate tax-to-GDP ratio explains relatively little: two countries with identical tax burdens but very different revenue compositions can achieve very different distributional outcomes (Adam, Kammas, and Lapatinas, 2015; OECD, 2024). And two countries with identical domestic tax compositions can face very different competitive and administrative environments depending on who their fiscal neighbours are. It is this second dimension, fiscal-network position, that existing empirical work has not addressed.

This paper fills that gap. We represent each country's tax system as a revenue-share vector across 28 OECD tax positions, the most granular harmonised decomposition available from OECD Revenue Statistics, and compute the pairwise cosine similarity of those vectors for 29 European economies over 1995–2022. The resulting similarity matrix defines a European tax-similarity network; while the underlying compositions shift over the sample, the estimated network is summarised by its element-wise median over 1995–2022, the fixed structural interaction matrix appropriate to claims about long-run structural exposure. Two indicators summarise each country's position in this network: tax diversification (normalised Shannon entropy of a country's row in the similarity network, capturing how evenly a country's fiscal resemblance is distributed across its European peers) and fiscal-network centrality (eigenvector centrality, capturing how closely a country's tax structure resembles those of many others simultaneously). A dynamic network Durbin model estimated by the defactored instrumental-variable procedure of Cui, Sarafidis, and Yamagata (2023), implemented via `spxtivdfreg` (Kripfganz and Sarafidis, 2025), then asks whether disposable-income inequality is shaped by domestic tax-structure composition, by a country's position in the fiscal network, or by the structural characteristics of the fiscal neighbourhood.

The paper makes three distinct contributions, each of which requires the network approach. The first is to represent tax systems not as lists of ratios but as positions in a network of fiscal

resemblance. The Shannon entropy of a country's row in this similarity network captures how broadly and evenly its fiscal resemblance is distributed across European peers, independently of the aggregate tax burden, a property that no single-instrument ratio can capture. The second is to replace geographic proximity with fiscal similarity as the relevant interaction structure. Countries with similar tax structures face similar competitive exposures, similar enforcement environments, and similar compliance pressures, and these structural resemblances are what the cosine-similarity network measures. The third contribution is to separate outcome diffusion (whether inequality levels propagate directly across similar countries) from structural diffusion (whether fiscal-neighbourhood characteristics independently condition domestic inequality). This separation, embodied in the Durbin versus SAR distinction, is the paper's central empirical finding: fiscal similarity matters not because it transmits inequality but because it creates structural exposure that shapes the redistributive effectiveness of national tax design. These three contributions jointly address a question that standard non-network panel approaches cannot address directly. A fourth contribution, and the paper's central policy-relevant finding, follows from the first two: the two indicators, both computed on the same fiscal-similarity network, capture distinct dimensions of tax-system architecture with distinct policy implications. Tax diversification is a domestically actionable margin, a government can broaden its revenue portfolio through national reform, whereas fiscal-network centrality is a structural exposure that no country can reduce unilaterally, because it is a property of the entire network. This asymmetry between a nationally actionable margin and a structurally determined constraint organises the paper's results and its policy implications.

Three robust findings emerge. Tax diversification is consistently and negatively associated with disposable-income inequality across all specifications, both network definitions, all sample periods, and both mechanism checks. Fiscal-network centrality is consistently and positively associated with inequality, supporting the interpretation that countries embedded at the centre of structurally similar peer systems face stronger competitive pressures on redistributive instruments. VAT compliance gaps among fiscal neighbours are positively associated with domestic inequality in full-sample estimates but not in post-2005 subsamples or first differences, indicating a long-run structural contextual association. The SAR restriction is rejected in favour of the Durbin model, confirming that fiscal-neighbour characteristics matter beyond direct Gini diffusion.

The empirical logic can be read in non-technical terms as follows. Tax diversification asks whether a country's own revenue portfolio is broad or concentrated. Fiscal-network centrality asks whether that portfolio resembles many other European tax systems, and therefore whether the country is structurally exposed to peer-system pressures. The network Durbin model then separates three objects: persistence in domestic inequality, direct inequality diffusion from fiscally similar neighbours, and structural diffusion through fiscal-neighbour characteristics. The econometric machinery controls for persistence, simultaneity in the network lag, and latent common shocks, while the substantive question remains whether fiscal architecture matters for disposable-income inequality.

The paper proceeds as follows. Section 2 reviews the relevant literature, establishes why the European Single Market is the right setting to study fiscal-network effects on inequality, and develops the conceptual mechanism. Section 3 describes the data and network construction. Section 4 presents the econometric framework. Sections 5 and 6 report main results and robustness checks. Section 7 concludes with policy implications.

2. Literature and Conceptual Framework

2.1. Tax Design and Distributional Consequences

The relationship between taxation and inequality has a long theoretical pedigree. The canonical trade-off between equity and efficiency under an optimal non-linear income tax (Mirrlees, 1971; Saez, 2001) established that tax design choices carry distributional consequences that extend well beyond the aggregate tax burden. The Atkinson–Stiglitz (1976) theorem provides a useful benchmark: under separable preferences and an optimally designed progressive income tax, differential commodity taxation adds nothing distributionally. But separability fails empirically, and the real world features tax avoidance, evasion, and administrative constraints that the benchmark ignores. When these departures are incorporated, composition matters: a euro of revenue raised through progressive personal income taxation achieves different distributional effects than a euro raised through a flat-rate value-added tax or a payroll-based social security contribution (Joumard, Pisu, and Bloch, 2012; IMF, 2014).

The specific distributional properties of different tax instruments are now well documented. Personal income taxes are generally progressive, though their redistributive effect depends critically on rate schedule design, the taxation of capital gains, and deductibility provisions. Piketty, Saez, and Stantcheva (2014) show that top tax rate reductions in OECD countries are associated with increases in top income shares without corresponding gains in economic growth, consistent with a rent-seeking mechanism rather than a labour supply response. Social security contributions are typically regressive relative to income above the contribution ceiling, and payroll taxes are largely shifted to wages in competitive labour markets (Gruber, 1997; Arulampalam, Devereux, and Maffini, 2012). VAT is regressive relative to income, though its redistributive consequences depend on exemptions, reduced rates, and the design of transfer programmes (Banks, Blundell, and Lewbel, 1997; Maier and Ricci, 2024). Corporate income tax incidence is contested, but labour may bear a significant share through reduced wages when capital is mobile (Fuest, Peichl, and Siegloch, 2018). These instrument-specific incidence patterns are the micro-foundation of the diversification mechanism that drives the paper's main result. The IMF (2015) documents that rising inequality across advanced economies reflects in part the inadequacy of tax-transfer systems to counter pre-distribution forces, a finding that motivates examining both the level and the composition of fiscal effort.

The broader fiscal capacity argument for revenue diversification draws on Besley and Persson's (2014) insight that fiscal capacity, the administrative and political infrastructure needed to collect broad-based taxes effectively, is a precondition for redistribution. A narrow tax base concentrates fiscal dependence on a small number of instruments that may be regressive, cyclically volatile, or administratively vulnerable to avoidance and evasion (Slemrod and Yitzhaki, 2002). A broader portfolio can sustain redistributive spending across economic conditions and reduces the pressure to rely on any single instrument with adverse distributional properties. The argument is one of portfolio resilience. Each instrument reviewed above carries a specific distributional vulnerability: personal income tax to avoidance at the top, social security contributions to regressivity above the ceiling, VAT to regressivity relative to income, and corporate tax to mobility. A system that depends heavily on any one of them inherits that vulnerability, whereas a system that spreads fiscal effort broadly diversifies away from it. The diversification index operationalises this portfolio-resilience argument in fiscal-similarity space: it is the entropy of a country's row in the similarity network, high when its revenue composition resembles many

European peers broadly and evenly and low when its resemblance, and the underlying instrument mix that generates it, concentrates on a narrow cluster. It is useful to fix the conceptual distinction between the two network indicators at the outset. Tax diversification, measured by entropy, captures the domestic dimension of tax-system architecture, how broadly a country's revenue composition positions it across the European field; a country's similarity position is generated by its own revenue mix, so this dimension remains domestically shaped. Fiscal-network centrality, developed in Section 2.3, captures the external dimension, how that domestic architecture is positioned relative to peer systems in the European competitive environment. The two measures answer different questions and, as the results in Section 5 confirm, carry different policy implications.

VAT compliance gaps occupy a specific place in this picture. Even when the statutory VAT structure is well-designed, the redistributive effectiveness of the system depends on actual collection relative to theoretical liability. Keen and Smith (2006) document that VAT fraud and evasion are structurally distinct from other tax compliance failures because the invoice chain mechanism creates specific incentive problems at each stage of production. Pomeranz (2015) shows that third-party information significantly improves VAT compliance through a self-enforcement mechanism, and that compliance responds strongly to audit risk. These findings imply that VAT compliance is not simply a domestic administrative matter but is affected by the information environment, audit cooperation, and enforcement norms prevailing in the wider market.

2.2. Cross-Country Evidence on Tax Composition and Inequality

The empirical literature on tax composition and inequality has grown substantially over the past decade, though almost all contributions share a common limitation: countries are treated as independent units beyond common time effects. Adam, Kammass, and Lapatinas (2015) find significant composition effects on inequality across a broad panel, with different bases exhibiting systematically different distributional associations. Alves and Afonso (2019) focus specifically on consumption-oriented tax structures and find that the redistributive consequences of VAT-heavy systems depend critically on the design of accompanying transfer programmes, with adverse distributional effects most pronounced where transfers are poorly targeted. Iosifidi and Mylonidis (2017) show for OECD economies that relative effective burdens on capital, labour, and consumption are more informative than statutory categories. Ciminelli, Ernst, Merola, and Giuliodori (2019) demonstrate that the distributional consequences of fiscal consolidations depend on whether adjustment relies on direct or indirect tax changes. Eydam and Qualo (2024) confirm that composition effects on inequality persist after controlling for country and time fixed effects. Maier and Ricci (2024) document that consumption taxation's redistributive consequences within the EU have become more heterogeneous as tax-system designs have diverged. The recent heterogeneous-panel study of Glavaški et al. (2025) confirms that direct tax forms reduce inequality while VAT increases it in European OECD economies, with effects varying significantly across countries.

A common thread in this literature is that system-level properties of tax structures matter. What the existing literature does not address, however, is the cross-country fiscal-network dimension: whether the distributional consequences of domestic tax composition depend on who the fiscal neighbours are. That is where this paper departs. De Paula, Rasul, and Souza (2025) show that fiscal interaction matrices inferred from data differ materially from geographic proximity and

that using the correct network matters for inference about policy interaction; the present paper extends that logic to the European setting.

2.3. Fiscal Interdependence in the European Single Market

The European Single Market institutionalises fiscal interdependence through the four freedoms, goods, services, capital, and labour, and an extensive body of EU law that constrains national tax autonomy. Five specific channels make fiscal spillovers stronger and more structurally determined inside Europe than in a generic cross-country panel.

VAT coordination pressure. EU VAT directives establish common rules for the tax base, minimum rates, and cross-border transaction treatment. Countries with similar VAT structures face direct comparability pressure and have historically adjusted toward structural convergence. This creates a specific channel through which VAT compliance conditions in one country's fiscal neighbourhood can affect the domestic enforcement environment, VAT fraud routes tend to exploit structural similarities between adjacent systems, a phenomenon extensively documented by the European Commission's VAT Gap monitoring (European Commission, 2025a, 2025b).

Social security and labour mobility. EU regulations ensure portability of social security entitlements across member states, making SSC design directly comparable and creating incentives for workers to account for SSC structures when making location decisions. Countries with high SSC burdens may face systematic downward pressure when similar neighbours reduce contribution rates.

Capital mobility and EU anti-avoidance frameworks. EU law restricts discriminatory taxation of non-resident income and capital. The Anti-Tax Avoidance Directive (ATAD) and the DAC transparency rules create shared administrative constraints that make the tax environments of similar countries increasingly interconnected. Countries with similar CIT structures face stronger competitive dynamics as tax-planning strategies exploiting rate differentials are more easily available.

Real-economy integration and supply chains. Deep product and labour market integration means that tax-induced changes in relative costs propagate rapidly across borders. Countries with similar production and trade structures, which tend to have similar tax structures, are more exposed to each other's fiscal choices through real-economy channels.

Policy learning and emulation. Governments observe and respond to fiscal reforms in similar countries. Redoano (2014) shows that EU membership amplifies strategic tax interactions relative to non-EU cross-country settings. Devereux and Loretz (2013) confirm that corporate tax competition is empirically pronounced among European economies with similar fiscal structures. Emulation and competitive pressure both operate most strongly between countries that are similar enough for direct comparison, precisely the countries that are most similar in our cosine-similarity network. The broader theoretical and empirical literature on tax competition is surveyed by Genschel and Schwarz (2011).

Taken together, these channels imply that the right interaction structure for modelling fiscal spillovers in Europe is not geographic contiguity but fiscal similarity. De Paula, Rasul, and Souza (2025) demonstrate for US state fiscal networks that interaction matrices inferred from data differ materially from geographic proximity and that using the correct network matters for

inference about policy interaction. The fiscal-similarity network constructed in this paper provides the European equivalent of that approach.

To see what structural similarity means in practice, take the case of a finance minister weighing a cut to the top personal income tax rate. A minister in a highly central country, such as Finland, France, or Croatia, finds that most peer systems carry a broadly similar schedule, so any reform is immediately benchmarked against fifteen or twenty comparable systems by tax advisers, business federations, opposition politicians, and the financial press; deviation from the European fiscal template is visible and costly. A minister in a peripheral country, such as Luxembourg or Sweden, operates from an already idiosyncratic structure, where deviation is expected and comparability is weaker.

Centrality, in this reading, measures how much unilateral fiscal room a country retains: the more central, the less room, and the stronger the competitive erosion of redistributive capacity over time. This is the network operationalisation of the asymmetric tax-competition result. The canonical model of Zodrow and Mieszkowski (1986) shows that strategic interaction between governments over mobile bases leads to the underprovision of public goods, and Wilson (1999) surveys the conditions under which this pressure is strongest. Bucovetsky (1991) establishes that more distinctive countries face lower competitive exposure, while Keen and Konrad (2013) provide the comprehensive treatment of competition and coordination across three decades of theory. Eigenvector centrality captures precisely the degree of symmetry in a country's relevant fiscal environment that this literature identifies as the driver of competitive intensity.

2.4. Conceptual Mechanism

The paper's empirical design separates three linked mechanisms through which fiscal-network structure might condition disposable-income inequality.

The domestic composition channel. Tax systems differ not only in aggregate fiscal effort but in the distribution of revenue across instruments with different incidence, mobility, and enforceability. A broader, more balanced revenue composition reduces dependence on any single instrument that may be regressive, cyclically volatile, politically constrained, or administratively vulnerable, and it places a country in a more evenly distributed position across its European peers rather than clustered with a few structurally near-identical systems. This positioning is what the diversification index captures, as the normalised entropy of a country's row in the similarity network: it is high when a country's revenue composition resembles many peers broadly and evenly, and low when its resemblance concentrates on a narrow cluster. Both readings point the same way, because the composition that diversifies a country's instrument base is also the composition that distributes its fiscal resemblance more evenly across the European field. This channel implies a negative direct effect of tax diversification on inequality. The nonlinear mechanism check tests whether this relationship exhibits diminishing returns at high levels, as would be expected when administrative fragmentation costs start to offset the diversification benefits.

The fiscal-network position channel. Countries embedded at the centre of the fiscal-similarity network, those whose revenue structures simultaneously resemble those of many other systems, face stronger competitive and emulative pressures. Centrality captures a country's structural exposure to the competitive dynamics of the European tax environment: a highly central country may find unilateral deviation from the common fiscal template more costly, as divergence risks

tax competition, avoidance arbitrage, or political pressure from comparator countries. This exposure may weaken redistributive capacity over time, as competitive adjustments to progressive or mobile-base instruments are more constrained. The channel implies a positive association between fiscal-network centrality and inequality. The policy-learning channel (which would predict the opposite sign) would prevail only if centrally embedded countries disproportionately adopt best-practice fiscal designs, a possibility that the data can adjudicate.

The compliance-externality channel. VAT compliance gaps measure the difference between theoretical and actual VAT revenue. When fiscal neighbours tolerate high VAT compliance gaps, cross-border arbitrage routes become more accessible, enforcement norms within the fiscal neighbourhood weaken, and the redistributive effectiveness of the domestic tax system may be undermined even if domestic compliance is adequate. This is a fiscal externality mediated through the similarity network: it is not bilateral but systemic. The indirect impact decomposition of the Durbin model, through which fiscal neighbours' VAT gaps affect domestic inequality via the network Leontief inverse, is the decomposition that allows this channel to be examined empirically. The direct domestic VAT effect is expected to be near zero once neighbourhood conditions are controlled, because the evidence is consistent with the fiscal neighbourhood's compliance environment playing a primary role rather than domestic compliance performance per se.

These three mechanisms are not mutually exclusive. A country can simultaneously benefit from domestic diversification, face competitive exposure through network centrality, and bear compliance externalities through its fiscal neighbourhood's VAT environment. The Durbin model separates domestic (β) and contextual (γ) effects, allowing each mechanism to be identified while the others are controlled. The three channels differ, finally, in a way that matters for institutional design. The domestic composition channel is within the reach of national policy: a government can broaden its revenue portfolio through tax reform regardless of what its neighbours do. The network-position channel is not: centrality is a property of the entire fiscal-similarity network, and a country cannot become peripheral simply by reforming its own system without diverging so far from the European template that the competitive and political costs become prohibitive. The compliance-externality channel shares this structural character, since it too is mediated by the neighbourhood rather than set domestically. This asymmetry, one nationally actionable margin and two structural exposures requiring coordinated response, maps directly onto the policy implications developed in Section 7.

2.5. Hypotheses

H1, Persistence and outcome diffusion. Disposable-income inequality is expected to be highly persistent. If fiscal similarity transmits distributional outcomes directly, the spatial lag of the Gini coefficient should be positive and significant. If not, the SAR restriction should be rejected in favour of the Durbin specification, indicating structural rather than outcome-based diffusion. Under a dissimilarity-weighted network, in which neighbours are defined by structural difference rather than similarity, competitive displacement rather than convergence is the operative channel, so outcome diffusion is predicted to be negative (H1b).

H2, Tax diversification. Countries occupying a more diversified position in fiscal-similarity space, generated by a broader and more balanced revenue composition, are expected to have lower disposable-income inequality through the domestic composition channel.

H3, Fiscal-network centrality. Countries more centrally embedded in the fiscal-similarity network are expected to be associated with higher inequality through the competitive-erosion channel.

H4, Structural fiscal-neighbour effects. Fiscal-neighbour characteristics are expected to condition domestic inequality beyond direct outcome diffusion. Among contextual variables, fiscal neighbours' VAT compliance gaps are the theoretically motivated primary channel.

3. Data and Network Construction

3.1. Sample and Dependent Variable

The analysis covers 29 European economies over 1995–2022 ($N = 29$, $T = 28$, 812 observations). The sample includes EU member states and closely integrated non-EU economies, Iceland, Norway, and Switzerland, whose EEA membership or bilateral agreements expose them to the same fiscal interaction patterns as EU members. Cyprus is the one EU member excluded, owing to incomplete coverage in OECD Revenue Statistics over the sample window. The dependent variable is the Gini coefficient of equivalised disposable household income from the Standardized World Income Inequality Database (SWIID v9.9; Solt, 2025), ranging from 21.9 to 38.2 with mean 28.8. The SWIID is preferred for its harmonised coverage across the full 1995–2022 window; a post-2005 robustness check addresses the potential concern that SWIID harmonisation relies more heavily on multiple imputation for earlier years, when EU-SILC cross-national coverage was more limited. The window spans the 2008–2012 financial and sovereign-debt crisis, during which several countries experienced sharp within-country inequality movements; the common component of this shock is absorbed by year fixed effects and the defactoring step, while country-specific responses are captured through the unemployment and government-effectiveness controls.

3.2. Revenue-Share Vector and $K = 28$ Tax Positions

Tax revenue data are drawn from OECD Revenue Statistics (OECD, 2025a) classified using the OECD interpretative guide (OECD, 2025b). Revenue shares are computed across $K = 28$ tax positions in the detailed "with-SSC" decomposition. These comprise five income tax positions (R1110–R1300, with PIT and CIT separated), seven social security contribution positions (R2110–R2400), one payroll position (R3000), ten property and wealth positions (R4110–R4600, including other recurrent property taxes R4600), four consumption positions (R5110, R5120, R5130, R5300), and one residual (R6000). The complete mapping with OECD codes is given in Appendix A (Table A1). Each country-year's 28-position revenue-share vector defines its position in the cosine-similarity network S constructed in Section 3.3; tax diversification is the normalised Shannon entropy of country i 's row in S , computed as

$$H_{it} = - \left[\frac{1}{\ln(28)} \right] \sum_{j=1}^{28} s_{ijt} \ln(s_{ijt}), 0 \leq H_{it} \leq 1. H_{it} = -(1/\ln 28) \cdot \sum_j s_{ijt} \ln s_{ijt}, \quad 0 \leq H_{it} \leq 1 \quad (j = 1, \dots, 28 \text{ peer countries})$$

where the s_{ijt} are the row-normalised similarity weights linking country i to its 28 European peers, so the index is computed on country i 's row of the annual similarity network S rather than over the 28 tax positions directly, and it is invariant to the aggregate tax-to-GDP ratio. The two objects are related but distinct: a country's similarity profile is generated by its own revenue

composition, so a more broadly balanced revenue mix tends to produce a more evenly distributed similarity profile, but the variable that enters the model and reproduces the panel mean in Table 1 is the similarity-space entropy, not the entropy of the 28 shares directly (the two are positively but only moderately correlated; see Appendix F.2). In concrete terms, the measure has a direct and clear reading. A country near the top of the diversification distribution resembles many European peers broadly and evenly, occupying a non-idiosyncratic but widely distributed position in fiscal-similarity space; a country near the bottom concentrates its fiscal resemblance toward a small cluster of peers. Because a country's position in similarity space is driven by its own revenue composition, diversification remains a domestically shaped margin: a government that broadens and rebalances its revenue mix moves toward a more diversified similarity position, which is the sense in which the measure is nationally actionable while centrality is not. For estimation, the diversification and centrality indices are min-max rescaled to the interval $[1.00001, 100]$ ($x = 1.00001 + 98.99999 (\text{raw} - \min)/(\max - \min)$), so each coefficient is read per point of the 0-100 index.

Three further controls enter the model: log GDP per capita at PPP (lypc , World Bank WDI), the ILO harmonised unemployment rate (u_ilo), and the World Governance Indicators Government Effectiveness index rescaled to $[0, 100]$ (ge , World Bank WGI). The lagged VAT compliance gap as a share of VAT Theoretical Tax Liability (vatgaplag , European Commission annual VAT Gap reports) is the primary contextual channel variable.

3.3. Fiscal-Similarity Network and Visualisation

The cosine similarity between countries i and j is $\text{sim}_{ij,t} = (s_{it} \cdot s_{jt}) / (\|s_{it}\| \|s_{jt}\|)$, bounded between 0 and 1. The similarity network W takes the element-wise median across annual matrices over 1995–2022, applies row-wise 95th-percentile pruning, and row-normalises so that $\text{sim}_{ij,t} =$

$$\frac{s_{it} \cdot s_{jt}}{\|s_{it}\| \|s_{jt}\|}, 0 \leq \text{sim}_{ij,t} \leq 1$$

$$\sum_{j=1}^N w_{ij} = 1, i = 1, \dots, N \quad w_{ii} = 0, i = 1, \dots, N \quad \sum_j w_{ij} = 1 \quad \text{for each } i, \quad \text{with } w_{ii} = 0, i = 1, \dots, N$$

yielding approximately two fiscal neighbours per country on average. The dissimilarity network W_{rev} weights by one minus cosine similarity and is used as a robustness check. Two points of construction matter for interpretation. The row-normalised, pruned matrix W is the interaction matrix used in estimation; pruning at the 95th percentile of the 28 possible links retains on average between one and two neighbours per country, and the matrix is directed before row-normalisation. Eigenvector centrality, by contrast, is computed on the full 29-by-29 median similarity matrix rather than on the pruned W , so that it measures a country's global position in the European fiscal structure rather than its local estimation neighbourhood. The symmetrised backbone in Figure 1(b), with 45 undirected edges, visualises that full structure and is not the estimation matrix. More generally, S and W play two analytically distinct roles, and it is worth stating this explicitly before the econometric framework. S is the information object: the node statistics computed on the full annual similarity matrices, the Shannon entropy of each country's row and its eigenvector centrality, describe each country's structural position in the European fiscal landscape and serve as the paper's two main regressors. W is the topology object: the pruned, row-normalised matrix specifies only which countries are fiscal neighbours for the spatial lag and contextual terms in estimation. Both derive from the same cosine-similarity

construction, but they serve different roles, and computing a node statistic on the pruned W rather than on the full S would be both technically incorrect and economically uninterpretable. The full-sample median is used for W because fiscal similarity among advanced European economies is structurally stable over the period: across the 28 annual networks eigenvector centrality varies within countries with a median coefficient of variation of 1.8% and a year-on-year rank correlation of 0.96, and the similarity-row entropy varies by under 0.1% within countries, so a time-invariant W avoids the complications of time-varying spatial weights at little cost. Section 6 nonetheless re-estimates the model on subperiod networks built separately from the first and second halves of the sample. Measuring proximity by the cosine similarity of share vectors follows the vector-space tradition of the technological-proximity literature (Jaffe, 1986), with revenue-category shares in place of patent-class shares.

Figure 1(a) displays the full 29×29 cosine-similarity matrix with countries ordered by Ward hierarchical clustering. A Central-Alpine cluster (Austria, Croatia, Finland, Switzerland) achieves within-group similarities above 0.80; a Western bloc (France, Italy, Spain, Belgium) forms a second cluster; and a Baltic cluster (Latvia, Lithuania, Netherlands) is identifiable. The key implication is visible in the off-diagonal cells: fiscal similarity does not map cleanly onto geographic proximity, Croatia and Switzerland are more fiscally similar than many geographically adjacent pairs, which is precisely why cosine-similarity weights capture the relevant interaction structure better than contiguity weights. Figure 1(b) shows the backbone network in Kamada–Kawai layout: Finland (EIG = 1.000) and Switzerland (EIG = 0.997) are the most central nodes; Luxembourg (EIG = 0.263) and Sweden (0.338) are the most peripheral.




Figure 1(a). Tax-structure similarity matrix (29 economies, 1995–2022 median). Ward hierarchical ordering; labels coloured by geographic region. Darker cells indicate higher cosine similarity. The block-diagonal structure reveals regional fiscal clustering, but the off-diagonal patterns confirm that fiscal proximity does not coincide with geographic proximity. Source: OECD Revenue Statistics, authors' calculations. The dark eastern block (Greece, Malta, Hungary, Croatia, Bulgaria, Romania) reflects structurally similar revenue profiles with high within-group cosine similarity, while Portugal and Spain stand out as a low-similarity column.




Figure 1(b). European tax-similarity network backbone (row-wise 95th-percentile pruning; 29 economies, 45 undirected edges). Node size = eigenvector centrality; node colour = geographic region; edge thickness = cosine similarity. Kamada–Kawai layout. Finland and Switzerland are co-central hubs; Luxembourg and Sweden are fiscal outliers. Source: OECD Revenue Statistics, authors' calculations.

3.4. Pre-Estimation Diagnostics

Table 1 reports descriptive statistics and Table 2 reports cross-sectional dependence (CD) diagnostics (Pesaran, 2004) and unit-root classifications (Pesaran, 2007; see Appendix C for full CIPS/CADF tables). Strong cross-sectional dependence is present in most raw variables, especially the Gini, tax diversification, eigenvector centrality, log GDP per capita, and unemployment. The VAT compliance gap is exceptional: its average cross-sectional correlation is negative ($CD = -2.21$, $p < 0.05$), pointing to asymmetric rather than common co-movement in VAT compliance dynamics across European fiscal systems. This negative co-movement is itself a substantive diagnostic and not merely a technical curiosity: when compliance improves in some countries it tends to deteriorate in others, an asymmetric pattern consistent with the

asymmetric VAT-compliance dynamics that motivate treating VAT gaps as network-dependent (Section 2.4). VAT enforcement dynamics appear competitive or substitutive across the European fiscal neighbourhood rather than commonly driven, a pattern consistent with the interdependence the network model is designed to capture. After two-way fixed-effects absorption, the compact TWFE residual retains $CD = -3.49$ ($p < 0.001$), confirming that latent common shocks persist beyond country and time effects and validating the defactored IV estimator rather than a simpler TWFE specification.

The CIPS/CADF tests reveal mixed integration orders: tax diversification rejects the unit-root null predominantly in levels; the VAT compliance gap and government effectiveness fail to reject in levels but strongly reject in first differences; and the Gini is highly persistent. Given this mixed picture, the preferred specification is a dynamic levels model interpreted as a stable dynamic network model rather than a panel cointegration model. Long-run steady-state multipliers are computed by `spxtivdfreg` via `estat impact, lr`; they are not cointegrating parameters. First-difference robustness checks for the persistent variables are reported in Section 6.

Table 1. Descriptive statistics (29 countries $\times$ 28 years, $N = 812$)

Variable	N	Mean	SD	Min	P25	Median	P75	Max	Within ratio
Gini (disp.)	812	28.84	3.43	21.9	26.0	28.3	31.6	38.2	0.322
<i>Tax diversification</i>	812	80.04	14.66	1.00	77.75	83.85	89.18	100.0	0.336
<i>Fiscal-network centrality</i>	812	75.47	18.09	1.00	66.53	78.68	88.71	100.0	0.370
<i>VAT gap (lagged, %)</i>	812	14.71	8.89	0.30	7.40	13.01	20.90	43.9	0.713
<i>Log GDP pc (PPP)</i>	812	10.11	0.775	8.17	9.53	10.19	10.70	11.63	0.154†
<i>Unemployment (%)</i>	812	8.23	4.36	1.81	5.04	7.14	10.23	27.69	0.592
<i>Gov. effectiveness</i>	812	82.16	13.91	33.33	72.46	83.33	94.20	100.0	0.273

Notes: Tax diversification = normalised Shannon entropy of country i 's row in the similarity network S over its 28 peer countries, min-max rescaled to $[1.00001, 100]$; Fiscal-network centrality = eigenvector centrality on S , min-max rescaled to $[1.00001, 100]$; VAT gap = VAT compliance gap as % of VTTL (European Commission), lagged one period; Gov. effectiveness = WGI Government Effectiveness, rescaled 0–100. Within ratio = SD after two-way FE / overall SD. † Log GDP per capita has limited within variation after FE absorption. Source: SWIID v9.9, OECD Revenue Statistics, World Bank WDI/WGI, European Commission VAT Gap Reports.

Table 2. Cross-sectional dependence and unit-root diagnostics

Variable	CD (raw)	p-val	Avg. corr.	CD (FE res.)	p-val	Integration class.
Gini (disp.)	19.63***	0.000	0.184	−3.52***	0.000	Highly persistent / near-I(1)
<i>Tax diversification</i>	24.49***	0.000	0.230	−1.05	0.294	Predominantly I(0)
<i>Fiscal centrality</i>	7.28***	0.000	0.068	−1.60	0.109	Borderline / I(1)-like
<i>VAT gap (lagged)</i>	−2.21**	0.027	−0.021	−2.85**	0.004	I(1); differences stationary
<i>Log GDP pc</i>	91.42***	0.000	0.857	−0.41	0.679	Predominantly I(0), borderline
<i>Unemployment</i>	27.06***	0.000	0.254	−0.49	0.623	Borderline I(0)
<i>Gov. effectiveness</i>	−0.48	0.628	−0.005	−2.81**	0.005	I(1); differences stationary
TWFE residual	−3.49***	0.000	−0.033			Latent factors remain after TWFE

Notes: CD = Pesaran (2004) cross-sectional dependence statistic. CD (FE res.) = CD test on two-way-FE residuals. CIPS/CADF (Pesaran, 2007), lags 0–3, intercept and trend. *** $p < 0.01$, ** $p < 0.05$. Full CIPS/CADF tables in Appendix C.

4. Econometric Framework

4.1. Dynamic Network Durbin Model

The maintained specification is the dynamic network Durbin model with multifactor interactive effects:

$$y_{it} = \alpha_i + \tau_t + \varphi y_{i,t-1} + \psi(Wy_t)_i + X'_{it}\beta + (WX_t)'_i\gamma + \lambda'_i f_t + \varepsilon_{it} \quad (1)$$

$$y_{i,t-1}\varphi(Wy_t)_i = \sum_{j=1}^N w_{ij}y_{jt} X'_{it}(WX_t)_i = \sum_{j=1}^N w_{ij}X_{jt} \lambda'_i f_t H_0: \gamma = 0$$

Country fixed effects α_i absorb time-invariant heterogeneity; a two-way within transformation eliminates both. The lagged Gini $y_{i,t-1}$ captures temporal persistence, with coefficient φ . The network lag $\sum_j w_{ij} y_{jt}$ tests outcome diffusion (H1), with coefficient ψ . The domestic regressor vector X_{it} contains tax diversification, fiscal-network centrality, the lagged VAT compliance gap, log GDP per capita, unemployment, and government effectiveness. The Durbin component $(\sum_j w_{ij} X_{jt})$ contains network-weighted values of the same fiscal variables for fiscal neighbours, testing structural diffusion (H4). The multifactor term $\lambda'_i f_t$ captures European-wide shocks that affect all countries simultaneously but with heterogeneous intensities (Sarafidis and Wansbeek,

2012). The NDM nests the dynamic SAR as $H_0: \gamma = 0$; rejection by a joint Wald test is the model-selection criterion.

4.2. Identification and Estimation

Three sources of endogeneity are addressed jointly. The contemporaneous network lag is endogenous because Gini outcomes are jointly determined within the fiscal network (Elhorst, 2014). The lagged dependent variable is endogenous in the presence of country-specific effects. Regressors may be correlated with latent common factors, inducing omitted-variable bias if factors are uncontrolled. The defactored IV procedure of Cui, Sarafidis, and Yamagata (2023) addresses all three jointly. In a first step, the common factors are projected out of the exogenous covariates by principal components. In a second step, the slope parameters, including the coefficient on the endogenous network lag, are estimated by instrumental variables, using the defactored covariates and their spatial counterparts as instruments. The model is then defactored a second time, using factors extracted from the first-step residuals, to obtain the final estimator. The estimator is asymptotically unbiased and requires no bias correction, unlike quasi-maximum-likelihood approaches subject to the incidental-parameters problem. The estimator is implemented in Stata via `spxtivdfreg` (Kripfganz and Sarafidis, 2025), an extension of `xtivdfreg` (Kripfganz and Sarafidis, 2021) to spatial settings. All specifications use 29 instruments. The number of common factors is selected by the Bai and Ng (2002) criterion (`factmax = 2`). Tax diversification and fiscal-network centrality are treated as predetermined, given the slow-moving political and legislative process of tax reform, an inertia well documented in the political economy of fiscal reform (Streeck and Thelen, 2005); the defactoring step removes common-factor contamination of this assumption. The Durbin regressors `WX`, formed from fiscal neighbours' characteristics, are treated as predetermined on the same grounds, since they are weighted averages of the same slow-moving structural variables. Both indices are computed on the full similarity network `S` rather than on the pruned interaction matrix `W`, so the regressors are not mechanically determined by the weight matrix used in estimation. And although both derive from the same network, their within-country partial correlation after two-way fixed-effects absorption is 0.41, confirming that they capture distinct dimensions of tax-system architecture over time rather than a single underlying factor: the entropy measures the breadth of a country's fiscal resemblance across peers, while eigenvector centrality measures its recursive embeddedness among well-connected peers.

The estimates reported below should be interpreted as robust conditional associations within a dynamic network model that controls for simultaneity, dynamic-panel bias, and latent common shocks. In the absence of external institutional instruments or natural experiments that would establish a quasi-experimental design, they do not constitute strict causal estimates of the effects of tax structure on inequality. The identification assumptions, the predeterminedness of the tax-structure variables and the approximate factor structure of the latent common shocks, are theoretically motivated and maintained rather than directly tested. The first-difference and lagged-specification robustness checks in Section 6 provide indirect support for these assumptions.

4.3. Impact Decomposition

$$S_{k(\psi)} = (I_N - \psi W)^{(-1)(\beta_k I_N + \gamma_k W)} \beta_k D E_k = \frac{1}{N} \text{tr}(S_{k(\psi)}) S_k^{LR(\varphi, \psi)}$$

$$= ((1 - \varphi) I_N - \psi W)^{(-1)(\beta_k I_N + \gamma_k W)}$$

In the network Durbin model, the marginal effects of regressor X_k are governed by the impact matrix $S_k(\psi) = (I_n - \psi W)^{-1} (\beta_k I_n + \gamma_k W)$, not by the structural coefficient β_k alone. The main diagonal elements are direct (own-country) effects, given by the average diagonal of $S_k(\psi)$; the off-diagonal elements are indirect (spillover) effects (LeSage and Pace, 2009). Long-run impacts are computed via estat impact, lr and represent steady-state responses under dynamic stability, not cointegrating parameters. This decomposition is central to identifying the compliance-externality channel: it separates the domestic direct effect of own VAT gap from the indirect network effect of fiscal neighbours' VAT gaps.

5. Results and Economic Interpretation

5.1. Outcome Diffusion or Structural Exposure? The Network Channel

Table 3 reports coefficient estimates across four model tiers ($N = 754$). All specifications include country and year fixed effects and share 29 instruments. Hansen J-tests are comfortable throughout. The estimated proportion of regression error variance attributable to the common factor component ranges from 0.31 to 0.45, confirming that latent common shocks account for a substantial share of residual cross-sectional dependence beyond two-way fixed effects. Log GDP per capita is consistently insignificant across specifications, as expected given its limited within-country variation after fixed-effects absorption (within ratio 0.154); it controls for cross-sectional development differences rather than within-country time variation.

The central model-selection result deserves careful interpretation because it reframes what the paper is about. In column (1), the SAR benchmark, the outcome diffusion coefficient, the response of domestic inequality to the network-weighted average Gini of fiscal neighbours, is positive but statistically insignificant. This is a meaningful null result, not a failure. It says that once domestic tax structure, macroeconomic controls, country and time fixed effects, and latent common factors are included, fiscally similar countries do not mechanically transmit their inequality levels to each other. The European co-movement in Gini coefficients visible in the raw data is explained by common shocks and time effects, not by fiscal-network contagion. The joint Wald test on the Durbin block in column (2) rejects the SAR restriction ($p = 0.018$). This shifts the interpretation from outcome diffusion to structural exposure: what matters is not the level of inequality among fiscal peers but the fiscal structures, competitive environments, and compliance conditions they represent. Column (2) is the preferred specification.

Table 3. Coefficient estimates, dynamic network Durbin model (defactored IV, spxtivdfreg)

	(1)	(2)	(3)	(4)
	SAR, W_sim	SDM core, W_sim	SDM sel., W_sim	SDM core, W_rev
L.Gini (φ)	0.852***	0.832***	0.830***	0.663***

	(0.020)	(0.038)	(0.031)	(0.030)
<i>Tax diversification</i>	−0.0137***	−0.0137***	−0.0139***	−0.0154***
	(0.0009)	(0.0017)	(0.0018)	(0.0017)
<i>Fiscal-network centrality</i>	0.0114***	0.0124***	0.0123***	0.0104***
	(0.0014)	(0.0019)	(0.0018)	(0.0017)
<i>VAT gap (lagged)</i>	0.0004	0.0003	0.0004	−0.0031**
	(0.0011)	(0.0015)	(0.0015)	(0.0013)
<i>Unemployment</i>	0.0125**	0.0151**	0.0157**	0.0302***
	(0.0055)	(0.0059)	(0.0058)	(0.0051)
<i>Log GDP per capita</i>	−0.361	−0.285	−0.284	−0.484
	(0.418)	(0.526)	(0.522)	(0.344)
<i>Gov. effectiveness</i>	−0.0009	−0.0013	−0.0012	0.0025
	(0.0016)	(0.0023)	(0.0023)	(0.0032)
W·Gini (outcome diffusion, ψ)	0.053	0.052	0.042	−0.077**
	(0.039)	(0.051)	(0.046)	(0.030)
<i>W·Tax diversification</i>	—	−0.0027	−0.0038	0.0006
		(0.0024)	(0.0024)	(0.0024)
<i>W·Fiscal centrality</i>	—	−0.0011	—	−0.0015
		(0.0023)		(0.0033)
<i>W·VAT gap (compliance ext.)</i>	—	0.0042*	0.0044**	0.0062**
		(0.0021)	(0.0019)	(0.0023)
Joint Durbin Wald test (p)	—	0.018	0.005	0.009
Hansen J (p)	0.460	0.181	0.229	0.479
Factor proportion ($\hat{\rho}_f$)	0.38	0.35	0.33	0.31
N	754	754	754	754

Notes: Robust standard errors in parentheses. All specifications include country and year fixed effects. Estimated by defactored IV (Cui, Sarafidis, and Yamagata, 2023) via `spxtivdfreg` (Kripfganz and Sarafidis, 2025). Columns (1)–(3): similarity network W_{sim} ; column (4): dissimilarity network W_{rev} . Tax diversification and fiscal-network

centrality are min-max rescaled to $[1.00001, 100]$ (see Section 3.2); mean-centred versions used only in mechanism specifications. ψ denotes the spatial autoregressive coefficient on W -Gini; ϕ denotes the autoregressive coefficient on the lagged dependent variable; $\hat{\rho}_f$ denotes the estimated proportion of total regression error variance attributable to the common factor component. 29 instruments (defactored exogenous covariates and their spatial lags). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5.2. Tax Diversification and Fiscal-Network Centrality

Tax diversification is the paper's most robust empirical finding. Across all four model tiers, both network definitions, all sample periods, and both mechanism checks, countries occupying a more broadly diversified position in fiscal-similarity space, resembling many European peers evenly rather than clustering with a few, are associated with lower disposable-income inequality. The coefficient stays negative and statistically significant across all nine specifications, with most estimates clustered between -0.011 and -0.017 , a degree of stability that is unusual in panel regression. The implication is substantive: where a country's revenue composition places it in the European fiscal field contains distributional information that the aggregate tax burden does not. Two countries raising 40% of GDP in taxes can achieve very different distributional outcomes if one's revenue composition concentrates its fiscal resemblance on a narrow cluster of near-identical systems while the other's distributes it broadly across the field.

The mechanism checks add texture. The positive quadratic term on tax diversification (significant at $p = 0.0002$) implies diminishing marginal returns to further diversification at high levels: the inequality-reducing effect of an additional unit is smaller when a country already occupies a broadly distributed position in similarity space. The interaction of tax diversification with government effectiveness (positive, $p = 0.0001$) shows that the inequality-reducing effect of diversification is attenuated in countries with stronger institutions, consistent with a substitution relationship: when administrative capacity is already high, the fiscal system can achieve redistribution from a more clustered similarity position. Full mechanism check results are in Appendix D.

Fiscal-network centrality is consistently and positively associated with inequality across all specifications. Countries whose revenue structures most closely resemble those of many other European systems simultaneously face higher disposable-income inequality. The direction and consistency of this finding support the competitive-erosion interpretation: being structurally similar to many fiscal peers amplifies exposure to tax competition, policy emulation, and avoidance arbitrage. The finding is consistent with Wilson's (1999) theoretical framework on tax competition and the underprovision of public goods, and with Devereux and Loretz's (2013) evidence that European corporate tax competition is empirically pronounced. The policy-learning channel does not dominate in this sample.




Figure 2. Tax diversification versus fiscal-network centrality (country means, 1995–2022).

Source: OECD Revenue Statistics, authors' calculations.

Read together, the two domestic coefficients carry an insight that neither carries on its own. Both are computed on the same fiscal-similarity network and enter the same specification, yet they have opposite signs and opposite policy logics. The negative diversification coefficient is a domestic margin: broadening the revenue portfolio reduces inequality, and the instrument for doing so is available to any government with sufficient administrative capacity and political will. The positive centrality coefficient is a structural constraint: embedding at the centre of the European fiscal template exposes a country to competitive dynamics that erode redistributive capacity, and relief from that constraint comes from coordinated European action rather than domestic reform. The asymmetry is not merely a statistical property of the estimates; it is a statement about the architecture of redistribution in an integrated market, what national tax reform can achieve at home is bounded by where a country sits in the fiscal network.

The case of Finland is instructive precisely because it appears paradoxical: it is the most central country in the network yet among Europe's most equal. The resolution is that the centrality coefficient captures a constraint, not an outcome. Finland faces the strongest competitive exposure in the European fiscal network, and the positive centrality–inequality association implies that this constrains its redistributive capacity; that Finland nonetheless achieves low inequality reflects the strength of its transfer system and wage-bargaining institutions, which the model captures only partially through government effectiveness and the macroeconomic controls. Finland is therefore not a counterexample to the finding but evidence for it: strong institutions can offset the competitive constraint that centrality identifies, and countries lacking Finland's institutional depth face the same constraint without the same compensating mechanisms.

5.3. The Compliance-Externality Channel: Impact Decomposition

The VAT compliance finding is structurally different from the two domestic findings and deserves a separate reading. Table 4 reports direct, indirect, and total impacts from the preferred column (2) specification. The own-country (direct) effect of the VAT compliance gap is near zero and insignificant: a country's own compliance performance does not independently associate with its inequality level once the neighbourhood is controlled. The indirect (network-mediated) effect, however, is large and statistically significant (long-run indirect effect +0.034, $p < 0.001$). This asymmetry is the paper's most novel finding from a methodological standpoint, and it is detectable only in the network Durbin model: a standard panel would find only the aggregate of direct and indirect effects and could not identify that the association is primarily contextual.

The finding suggests that what matters for domestic inequality is not primarily how compliant one's own VAT system is, but how compliant the fiscal neighbourhood is. When fiscal neighbours tolerate high VAT gaps, arbitrage routes become more accessible, audit cooperation weakens, and the redistributive effectiveness of the domestic tax system may be undermined through the network structure. This is consistent with the compliance-externality channel described in Section 2.4, and is supported in full-sample estimates. Because the association weakens after 2005 and in first differences, it is interpreted as a long-run structural contextual finding rather than a contemporaneous short-run transmission mechanism.

Table 4. Short-run and long-run impact decomposition, preferred specification, column (2)

	SR direct	SR indirect	SR total	LR direct	LR indirect	LR total
<i>Tax diversification</i>	−0.0138*** (0.0017)	−0.0035** (0.0017)	−0.0173*** (0.0017)	−0.0866*** (0.0207)	−0.0546** (0.0207)	−0.141*** (0.021)
<i>Fiscal centrality</i>	0.0124*** (0.0019)	−0.0005 (0.0019)	0.0119*** (0.0019)	0.0758** (0.0270)	0.0211 (0.0270)	0.0969*** (0.027)
<i>VAT gap (lagged)</i>	0.0004 (0.0015)	0.0044** (0.0015)	0.0048** (0.0015)	0.0050 (0.0096)	0.0338*** (0.0096)	0.0388*** (0.010)
<i>Unemployment</i>	0.0151** (0.0059)	0.0008 (0.0059)	0.0159** (0.0059)	0.0930** (0.0396)	0.0365 (0.0396)	0.130*** (0.040)

Notes: Delta-method standard errors in parentheses. SR = short-run; LR = long-run steady-state response computed via $\text{estat impact, lr (spxtivdfreg)}$. The impact matrix for regressor X_k is $Sk(\psi) = (IN - \psi W)^{-1}(\beta_k IN + \gamma_k W)$. Average direct effects: own-country responses. Average indirect effects: network-mediated spillovers. Note the near-zero direct effect and large significant indirect effect for the VAT compliance gap: this pattern is consistent with the compliance-externality interpretation, in which the fiscal neighbourhood's compliance environment plays a primary role. *** $p < 0.01$, ** $p < 0.05$.

The long-run magnitudes are economically meaningful. A 10-point increase in the tax-diversification index is associated with approximately 1.4 Gini points lower disposable-income inequality in the long run, and a 10-point increase in fiscal-network centrality with approximately 1.0 Gini point higher inequality. These effects are non-trivial relative to the observed cross-country dispersion of disposable-income inequality, though they are full-sample long-run associations rather than causal treatment effects for any single country.

6. Robustness

Table 5 summarises robustness checks across nine specifications. The two core domestic findings, tax diversification reduces inequality, fiscal-network centrality raises it, survive every check without exception. The VAT compliance finding requires more nuanced discussion.

Table 5. Robustness summary, domestic and contextual coefficients

Specification	Tax divers.	Fiscal cent.	W·VAT gap	Hansen J (p)
Baseline, column (2), W_sim	−0.0137***	0.0124***	+0.0042*	0.181
Column (3): selected Durbin (W·D, W·VAT)	−0.0139***	0.0123***	+0.0044**	0.229
Column (4): dissimilarity network W_rev	−0.0154***	0.0104***	+0.0062**	0.479
R1: full WX controls (add W·lypc, W·u, W·ge)	−0.0166***	0.0114***	+0.0071**	0.334
R2: post-2005 subsample (2005–2022)	−0.0147***	0.0102***	−0.0055 n.s.	0.926
R3: first-difference specification	−0.0109***	0.0055***	+0.0017 n.s.	0.588
R4: lagged tax structure (L.se1, L.eig1)	−0.0034**	0.0050***	+0.0046 n.s.	0.442
M5: nonlinear diversification (se1 ²)	−0.0140***	0.0123***	+0.0060**	0.366
M7: institutional moderation (se1 × ge)	−0.0173***	0.0104***	+0.0022 n.s.	0.626

*Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$; n.s. = not significant at 10%. Tax diversification and fiscal-network centrality are min-max rescaled to [1.00001, 100] (Δ -versions for R3). In R2 the joint Durbin Wald test fails to reject the SAR restriction ($p = 0.500$) and W·VAT gap reverses sign (negative, insignificant, $p = 0.151$), indicating the structural diffusion channel is absent or weaker post-2005. In R3 (first differences) the domestic effects survive but the VAT compliance spillover is not significant ($p = 0.427$), consistent with a long-run structural rather than short-run contemporaneous channel. In R4 the dynamic stability condition is violated; long-run impacts are not computed. M5 and M7 use mean-centred variables c_se1 , c_ge .*

The dissimilarity network result (column 4) is informative. Using W_{rev} preserves the domestic diversification and centrality findings exactly. The outcome diffusion coefficient turns negative and significant (-0.077 , $p = 0.011$), consistent with competitive displacement across structurally different systems. The positive $W \cdot VAT$ gap remains significant (0.0062 , $p = 0.008$). The persistence of this association under a network that defines fiscal neighbours in the opposite way to the baseline, by structural dissimilarity rather than similarity, suggests that the full-sample association is not mechanically tied to the baseline similarity matrix. However, because the similarity and dissimilarity networks represent different channels of fiscal exposure, this is interpreted as robustness of the sign rather than confirmation of the same theoretical mechanism. The contextual coefficient on $W \cdot Tax$ diversification is itself negative under the similarity network and positive under the dissimilarity network, consistent with fiscal neighbours' diversification operating through opposite competitive channels under the two network definitions; neither estimate reaches significance, however, and the robust result is the domestic diversification effect rather than its spatial counterpart.

Figure 3. Coefficient stability across specifications




Figure 3. Coefficient stability across the eleven robustness specifications. Markers are point estimates; bars are 95% confidence intervals. Source: authors' calculations.

A further battery confirms that the two domestic findings are not artefacts of the network's construction. Pruning the similarity matrix at the 90th, 85th, 80th, and 75th percentiles (matrices W_2 to W_5) leaves diversification negative and significant throughout (between -0.0144 and -0.0169) and centrality positive and significant (0.0105 to 0.0130), with Hansen J p-values between 0.21 and 0.69 . Splitting the sample into its first and second halves, each estimated on its own subperiod network (the early and late networks), preserves both findings (diversification -0.0183 first half, -0.0134 second half; centrality 0.0059 first half, 0.0123 second half). Replacing eigenvector centrality with PageRank or betweenness preserves the positive network-position result (PageRank 0.0120 , betweenness 0.0018). Adding total tax revenue as a share of GDP leaves the diversification and centrality coefficients intact, while tax-to-GDP itself is insignificant. This indicates that the diversification result is not a proxy for aggregate fiscal effort. A further governance check, using Control of Corruption in place of government effectiveness, confirms that the domestic findings are not driven by the baseline institutional

control. We deliberately do not condition on individual tax-base shares or statutory rates, since these are components of the very portfolio whose architecture the analysis measures: conditioning on them would change the estimand to a residual association net of the portfolio's own internal channels, which is why the tax-to-GDP check is the appropriate fiscal-capacity test. Figure 3 summarises the eleven specifications: the diversification coefficient is negative in every one and the centrality coefficient positive in every one. The smaller betweenness coefficient is expected rather than inconsistent: betweenness captures a country's role as a bridge between otherwise separated clusters, a different structural property from eigenvector centrality's overall embeddedness, so the two are not equivalent substitutes.

CIPS/CADF tests classify the VAT compliance gap and government effectiveness as I(1)-like, motivating the first-difference robustness check (R3). In first differences, the domestic effects of tax diversification and fiscal-network centrality survive unambiguously, confirming that the levels estimates are not spurious artefacts of persistence. The VAT compliance spillover is not significant in first differences, nor in the post-2005 subsample. Combined with the post-2005 structural instability (the SAR restriction is not rejected in R2), this points to a long-run association concentrated in the earlier part of the sample rather than a year-to-year transmission mechanism. The preferred specification is therefore maintained as a stable dynamic levels model and the VAT compliance result is interpreted accordingly.

A final robustness check addresses the measurement of tax diversification directly. Because the diversification index is computed on the similarity network S rather than on the revenue shares themselves (Section 3.2 and Appendix F), we re-estimate the preferred specification (Table 6) replacing it with a direct Shannon entropy index computed from the 28 OECD tax-position revenue shares within each country (H_share , rescaled to the same $[1.00001, 100]$ interval). The two indices are only moderately correlated ($r = 0.46$), confirming that they capture related but non-identical dimensions of fiscal structure. The direct share-entropy coefficient is -0.0033 and remains negative and statistically significant ($p = 0.015$), while fiscal-network centrality (0.0086) and the fiscal-neighbour VAT-gap spillover (0.0082) retain their signs and significance. The inequality-reducing association of tax diversification is therefore not an artefact of the network-based construction: it holds whether diversification is measured as a country's position in fiscal-similarity space or as the breadth of its revenue composition across the 28 tax positions directly.

	(1) Baseline ($se1$)	(2) Direct share entropy (H_share)
Tax diversification	-0.0137^{***} (0.0017)	-0.0033^{**} (0.0014)
Fiscal-network centrality	0.0124^{***} (0.0019)	0.0086^{***} (0.0015)
W·VAT gap	0.0042^* (0.0021)	0.0082^{***} (0.0030)
Hansen J (p)	0.181	0.155
N	754	754

Table 6. Tax-diversification measure robustness: preferred specification (Table 3, column 2) with the network-based similarity-row entropy ($se1$) versus a direct 28-position revenue-share entropy (H_share). Both are rescaled to $[1.00001, 100]$. Robust standard errors in parentheses. $*** p < 0.01$, $** p < 0.05$.

7. Conclusion and Policy Implications

7.1. Summary of Findings

This paper started from a puzzle: Finland and Romania participate in the same Single Market and face the same EU fiscal directives, yet maintain a persistent disposable-income inequality gap that standard domestic-tax explanations account for only in part. The network lens offers a structural complement: in the European Single Market, national tax systems are not independent choices but are made inside a fiscal neighbourhood of structurally similar systems, connected through capital mobility, EU directives, shared administrative constraints, and competitive dynamics. Finland's deep embedding in the European fiscal-similarity network constrains its redistributive manoeuvre; Romania's peripheral position does not. We formalised this observation by representing each country's tax system as a revenue-share vector across 28 OECD positions and constructing a cosine-similarity network from those vectors. Using a dynamic network Durbin model with defactored IV, we then addressed a question that standard non-network panel approaches cannot reach directly: does disposable-income inequality in Europe depend not only on domestic tax-system composition, but on fiscal-network position and fiscal-neighbourhood characteristics? The estimates throughout should be read as robust dynamic-network associations rather than quasi-experimental causal effects.

The answer is yes, with important qualifications. Tax diversification is the paper's most robust finding: a more diversified position in fiscal-similarity space, reflecting a broader and more balanced revenue composition, is consistently associated with lower inequality across all specifications, sample periods, and network definitions, with a long-run total effect of approximately -0.14 . In practical terms, a country moving from the 25th to the 75th percentile of the diversification distribution, a difference of roughly 11 points on the 0 to 100 index, is associated with a long-run reduction in the disposable-income Gini of about 1.6 points. Fiscal-network centrality is also robust and positive ($+0.097$): countries more deeply embedded in the European fiscal-similarity network are associated with higher inequality, consistent with competitive pressure on redistributive instruments. The SAR restriction is rejected in favour of the Durbin specification, establishing that fiscal-neighbour characteristics matter beyond direct Gini transmission. The VAT compliance gap of fiscal neighbours is positively associated with domestic inequality in full-sample estimates, primarily through the indirect network channel, though this association weakens after 2005 and disappears in first differences. These findings hang on one organising insight. Fiscal-system architecture has two dimensions computed from the same fiscal-similarity network: a domestic dimension, the breadth of a country's fiscal resemblance across peers, that national policy can act on through its own revenue composition, and a network dimension, structural exposure, that it cannot. The distinction between them is the distinction between what tax reform can achieve at home and what it can achieve only in concert with neighbours.

7.2. Policy Implications and Boundaries

The policy implications are structured by the evidence hierarchy, with the robustness of each calibrated to the robustness of its empirical basis, and by the domestic-versus-structural asymmetry identified above. That asymmetry gives the implications an institutional logic. The diversification finding belongs to national fiscal strategy, because a country's diversification position is one a government can shift unilaterally through its own revenue composition. The

centrality finding belongs to European coordination frameworks, because structural exposure cannot be reduced by domestic reform alone. The more tentative compliance finding belongs to a third level, shared enforcement infrastructure, that bridges the two.

The strongest implication concerns domestic revenue-base design. The consistent negative association between tax diversification and inequality suggests that distributional policy should not focus solely on aggregate tax ratios or on individual instrument rates. Where a country's revenue composition places it in the European fiscal field, broadly distributed across peers rather than clustered with a few, appears to carry independent distributional information. This is consistent with Besley and Persson's (2014) fiscal-capacity argument: administrative and political investment in a broad tax base is a prerequisite for sustained redistribution. The implication is one of targeting rather than a generic call to diversify. The economies in the most clustered positions in fiscal-similarity space, whose concentrated revenue compositions place them close to a few near-identical peers rather than spread across the European field, face the steepest marginal inequality-reducing gains from rebalancing their revenue composition toward a more broadly distributed position. Because that position is generated by a country's own revenue mix, this remains a domestically actionable margin, and European Semester recommendations on revenue-base diversification are therefore most valuable where a country's fiscal resemblance is most narrowly concentrated. The nonlinear mechanism result adds a practical calibration: the marginal inequality-reducing effect of further diversification is smaller for countries already in a broadly distributed position, suggesting that beyond a threshold, composition changes pursued without the administrative capacity to support them add fragmentation costs without additional redistributive benefit.

The second implication concerns fiscal-network centrality and the limits of unilateral redistribution. The positive centrality coefficient is consistent with the idea that countries structurally similar to many peers face stronger competitive constraints on maintaining progressive rate schedules, particularly for mobile bases. This provides distributional motivation for supranational coordination on the most mobile tax bases, such as corporate income, capital gains, and financial transactions, complementing the revenue-efficiency case established in the tax-competition literature (Wilson, 1999; Devereux and Loretz, 2013; Sørensen, 2004, on the welfare gains from partial coordination in an EU-like setting). Because the centrality measure is built from the full revenue portfolio rather than corporate tax alone, the coordination logic extends beyond corporate income tax to other mobile instruments. The most central economies, Finland and Switzerland and the others clustered near the centre, are precisely those for which unilateral redistribution faces the strongest structural constraint, and for which EU-level minimum standards deliver the largest marginal distributional relief. The most peripheral economies, such as Luxembourg and Sweden, retain more fiscal room but may gain less from coordination. The ongoing implementation of OECD Pillar Two and the EU minimum corporate tax directive is directly relevant: minimum standards limit the competitive externalities associated with fiscal-network centrality, protecting the redistributive capacity of member states without requiring full harmonisation. The proposed Business in Europe Framework for Income Taxation (BEFIT) extends this coordination logic to a common corporate-tax base; the centrality finding supplies a distributional rationale for such frameworks that complements the single-market and revenue-efficiency case the Commission has advanced.

The third implication concerns VAT enforcement coordination, and here the evidentiary basis is more qualified. The full-sample estimates show a positive and network-mediated association

between fiscal neighbours' VAT compliance gaps and domestic inequality. The mechanism, facilitated cross-border arbitrage, weakened enforcement norms, undermined redistributive effectiveness, is consistent with Keen and Smith's (2006) analysis of VAT fraud across integrated markets and with Pomeranz's (2015) evidence that VAT compliance depends on the information and enforcement environment of the system as a whole. However, because this association weakens after 2005 and disappears in first differences, it should not be interpreted as definitive proof that VAT enforcement coordination reduces inequality. The more cautious interpretation is that VAT enforcement coordination is a complementary distributional policy margin, it may help protect the redistributive capacity of national tax systems by reducing compliance externalities, but the channel appears to be structural and long-run rather than contemporaneous.

This evidence speaks directly to a live European policy debate. The VAT in the Digital Age (ViDA) package, formally adopted in March 2025 (Council Directive (EU) 2025/516), introduces real-time digital reporting requirements for intra-EU transactions, new platform economy VAT obligations, and expanded single-registration mechanisms. Its primary rationale is reducing documented VAT revenue losses across member states (European Commission, 2025a, 2025b). The paper adds a distributional rationale: reducing VAT compliance heterogeneity across the European fiscal neighbourhood may also reduce the compliance externalities that, in full-sample estimates, are associated with higher domestic inequality. The OECD (2024) Taxation and Inequality report, commissioned for G20 Finance Ministers, frames tax systems as capable of mitigating or exacerbating inequality depending on how reform scope is directed at the distribution of income and wealth; the paper's evidence supports extending that framing to the enforcement dimension of VAT systems in integrated markets.

Three inferences the paper does not support deserve explicit statement, since policy readers may otherwise over-interpret the findings. The paper does not establish that any named country should restructure its tax code in a particular direction. The associations are long-run and cross-national, not a treatment protocol for a single economy, and entropy captures the balance of the portfolio, not the direction in which any one of the 28 positions should be expanded or contracted. It does not show that peripheral countries achieve better redistribution through their idiosyncrasy: lower centrality reduces competitive exposure but does not guarantee better outcomes where institutional capacity is weak. And it does not advocate full tax harmonisation: the instruments the findings support are minimum standards that relax competitive pressure while preserving national fiscal autonomy, not harmonisation that eliminates it. These are the limits within which the implications above should be read.

The broader implication is conceptual and durable. Tax policy in Europe is not a purely national exercise. The distributional consequences of domestic tax design depend partly on the fiscal neighbourhood in which that design is made. European redistribution policy should therefore treat national tax systems as interdependent fiscal structures, recognising that domestic base diversification, supranational coordination of competitive tax bases, and coordinated VAT enforcement may operate as complementary rather than alternative policy margins.

7.3. Limitations and Future Research

Two kinds of limitation shape these findings and are worth distinguishing. The first are limitations of this paper proper. The sample covers advanced European economies, and the

similarity-space diversification mechanism may operate differently in lower-income settings where administrative capacity constraints are more severe. The network construction uses cosine similarity of revenue-share vectors; alternatives based on Aitchison compositional geometry, or on correlations between residuals after extracting fixed effects, are natural extensions. The VAT compliance gap post-2005 weakness invites subperiod analysis as longer time series become available. And the predeterminedness assumption for tax-structure variables, rather than full instrumentation by exogenous institutional changes, is maintained rather than tested; future work with suitable natural experiments could strengthen the identification. A second set of constraints is of a different character: these are frontier limitations of the estimator itself rather than choices made for convenience here. The spatial weight matrix is treated as a fixed structural object computed from observed revenue shares before estimation, which is the source of the paper's contribution rather than a simultaneity to be purged, and the defactoring step controls for the common-factor channel through which the network and the outcome could otherwise co-move. We nonetheless treat the network as exogenous to the inequality process, an assumption that would be tested most directly by exploiting exogenous shocks to fiscal similarity such as EU accession waves or directive-driven harmonisation events. The network is held fixed at its 1995–2022 element-wise median. Estimating a genuinely time-varying weight matrix, and recovering the heterogeneous-slope analogue of the impact decomposition following Chen, Cui, Sarafidis, and Yamagata (2025) and Kripfganz and Sarafidis (2026), awaits inferential machinery that the current implementation does not yet deliver. The mean-group estimator does not presently return the direct, indirect, and total impact matrix that is this paper's central analytical output. These are boundaries of the methodological frontier, shared by all applications of this class of estimator, rather than weaknesses specific to the present study.

Declarations

Funding: The authors received no specific funding for this work.

Conflicts of interest: The authors declare no conflict of interest.

Data availability: The data can be available from the authors upon request.

Author contributions: Achillefs Karakotsios, Thomas Panagiotou, and Constantinos Katrakilidis contributed to the conceptualization, methodological design, data construction, analysis, interpretation of results, and writing of the manuscript. All authors reviewed and approved the final version.

References

- Adam, A., Kammas, P., and Lapatinas, A. (2015). Income inequality and the tax structure: Evidence from developed and developing countries. *Journal of Comparative Economics*, 43(1), 138–154.
- Alves, J., and Afonso, A. (2019). Tax structure for consumption and income inequality: An empirical assessment. *SERIEs*, 10(3–4), 337–364.
- Arulampalam, W., Devereux, M. P., and Maffini, G. (2012). The direct incidence of corporate income tax on wages. *European Economic Review*, 56(6), 1038–1054.
- Atkinson, A. B., and Stiglitz, J. E. (1976). The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics*, 6(1–2), 55–75.
- Bai, J., and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1), 191–221.
- Banks, J., Blundell, R., and Lewbel, A. (1997). Quadratic Engel curves and consumer demand. *Review of Economics and Statistics*, 79(4), 527–539.
- Besley, T., and Persson, T. (2014). Why do developing countries tax so little? *Journal of Economic Perspectives*, 28(4), 99–120.
- Bucovetsky, S. (1991). Asymmetric tax competition. *Journal of Urban Economics*, 30(2), 167–181.
- Chen, J., Cui, G., Sarafidis, V., and Yamagata, T. (2025). IV estimation of heterogeneous spatial dynamic panel models with interactive effects. Working Paper. arXiv:2501.18467.
- Ciminelli, G., Ernst, E., Merola, R., and Giuliadori, M. (2019). The composition effects of tax-based consolidation on income inequality. *European Journal of Political Economy*, 57, 107–124.
- Cui, G., Sarafidis, V., and Yamagata, T. (2023). IV estimation of spatial dynamic panels with interactive effects: Large sample theory and an application on bank attitude toward risk. *The Econometrics Journal*, 26(2), 124–146.
- de Paula, Á., Rasul, I., and Souza, P. C. L. (2025). Identifying network ties from panel data: Theory and an application to tax competition. *The Review of Economic Studies*, 92(4), 2691–2729.
- Devereux, M. P., and Loretz, S. (2013). What do we know about corporate tax competition? *National Tax Journal*, 66(3), 745–774.
- Elhorst, J. P. (2014). *Spatial econometrics: From cross-sectional data to spatial panels*. Springer.
- European Commission. (2025a). VAT gap in Europe, Report 2025. Directorate-General for Taxation and Customs Union. Publications Office of the European Union.
- European Commission. (2025b). VAT in the Digital Age (ViDA), Council Directive (EU) 2025/516 of 11 March 2025. Official Journal of the European Union.

Eydam, U., and Qualo, H. (2024). Income inequality and taxes, An empirical assessment. *Applied Economics Letters*, 31(18), 1828–1835.

Fuest, C., Peichl, A., and Siegloch, S. (2018). Do higher corporate taxes reduce wages? Micro evidence from Germany. *American Economic Review*, 108(2), 393–418.

Genschel, P., and Schwarz, P. (2011). Tax competition: A literature review. *Socio-Economic Review*, 9(2), 339–370.

Glavaški, O., Beker Pucar, E., Beljić, M., and Stojkov, S. (2025). Rethinking tax systems: How heterogeneous tax mix shapes income inequality in European OECD economies. *Journal of Risk and Financial Management*, 18(5), 279.

Gruber, J. (1997). The incidence of payroll taxation: Evidence from Chile. *Journal of Labor Economics*, 15(S3), S72–S101.

International Monetary Fund. (2014). Fiscal policy and income inequality. IMF Policy Paper.

International Monetary Fund. (2015). Causes and consequences of income inequality: A global perspective. IMF Staff Discussion Note 15/13.

Iosifidi, M., and Mylonidis, N. (2017). Relative effective taxation and income inequality: Evidence from OECD countries. *Journal of European Social Policy*, 27(1), 57–76.

Jaffe, A. B. (1986). Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value. *American Economic Review*, 76(5), 984–1001.

Journard, I., Pisu, M., and Bloch, D. (2012). Tackling income inequality: The role of taxes and transfers. *OECD Journal: Economic Studies*, 2012(1), 37–70.

Keen, M., and Konrad, K. A. (2013). The theory of international tax competition and coordination. In A. J. Auerbach, R. Chetty, M. Feldstein, and E. Saez (Eds.), *Handbook of Public Economics*, Vol. 5, pp. 257–328. Elsevier.

Keen, M., and Smith, S. (2006). VAT fraud and evasion: What do we know and what can be done? *National Tax Journal*, 59(4), 861–887.

Kripfganz, S., and Sarafidis, V. (2021). Instrumental variable estimation of large-T panel data models with common factors. *The Stata Journal*, 21(3), 659–686.

Kripfganz, S., and Sarafidis, V. (2025). Estimating spatial dynamic panel data models with unobserved common factors in Stata. *Journal of Statistical Software*, 113(6), 1–27.

Kripfganz, S., and Sarafidis, V. (2026). Chasing opportunity: Spillovers and drivers of US state population growth. *Spatial Economic Analysis*, 21(1), 61–79.

LeSage, J. P., and Pace, R. K. (2009). *Introduction to spatial econometrics*. CRC Press.

Maier, S., and Ricci, M. (2024). The redistributive impact of consumption taxation in the EU. *Economic Analysis and Policy*, 81, 738–755.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2), 175–208.

- OECD. (2024). Taxation and inequality: OECD report to G20 Finance Ministers and Central Bank Governors. OECD Publishing.
- OECD. (2025a). Revenue Statistics 2025. OECD Publishing.
- OECD. (2025b). The OECD classification of taxes and interpretative guide. In Revenue Statistics 2025. OECD Publishing.
- Pesaran, M. H. (2004). General diagnostic tests for cross section dependence in panels. Cambridge Working Papers in Economics No. 0435.
- Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of Applied Econometrics*, 22(2), 265–312.
- Piketty, T., Saez, E., and Stantcheva, S. (2014). Optimal taxation of top labor incomes: A tale of three elasticities. *American Economic Journal: Economic Policy*, 6(1), 230–271.
- Pomeranz, D. (2015). No taxation without information: Deterrence and self-enforcement in the value added tax. *American Economic Review*, 105(8), 2539–2569.
- Redoano, M. (2014). Tax competition among European countries: Does the EU matter? *European Journal of Political Economy*, 34, 353–371.
- Saez, E. (2001). Using elasticities to derive optimal income tax rates. *The Review of Economic Studies*, 68(1), 205–229.
- Sarafidis, V., and Wansbeek, T. (2012). Cross-sectional dependence in panel data analysis. *Econometric Reviews*, 31(5), 483–531.
- Slemrod, J., and Yitzhaki, S. (2002). Tax avoidance, evasion, and administration. In A. J. Auerbach and M. Feldstein (Eds.), *Handbook of Public Economics*, Vol. 3, pp. 1423–1470. Elsevier.
- Solt, F. (2025). Standardized World Income Inequality Database (SWIID), Version 9.9. fsolt.org.
- Sørensen, P. B. (2004). International tax coordination: Regionalism versus globalism. *Journal of Public Economics*, 88(6), 1187–1214.
- Streeck, W., and Thelen, K. (2005). *Beyond Continuity: Institutional Change in Advanced Political Economies*. Oxford University Press.
- Wilson, J. D. (1999). Theories of tax competition. *National Tax Journal*, 52(2), 269–304.
- Zodrow, G. R., and Mieszkowski, P. (1986). Pigou, Tiebout, property taxation, and the underprovision of local public goods. *Journal of Urban Economics*, 19(3), 356–370.

Appendix A. Country Sample and Tax Code Mapping (K = 28)

The 29-country panel: Austria, Belgium, Bulgaria, Croatia, Czechia, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Norway, Poland, Portugal, Romania, Slovak Republic, Slovenia, Spain, Sweden, Switzerland. Iceland, Norway, and Switzerland participate via EEA membership or bilateral agreements.

Table A1. OECD Revenue Statistics codes mapped to K = 28 tax positions (with SSC)

Broad category	Position / subcategory	OECD code (Stata variable)
PIT & CIT (1000), 5 positions	Taxes on income & profits of individuals	<i>1110 (R1110)</i>
	Capital gains of individuals	<i>1120 (R1120)</i>
	Taxes on income & profits of corporations	<i>1210 (R1210)</i>
	Capital gains of corporations	<i>1220 (R1220)</i>
SSC (2000), 7 positions	Unallocable between individuals / corporations	<i>1300 (R1300)</i>
	SSC by employers, payroll basis	<i>2110 (R2110)</i>
	SSC by employers, income tax basis	<i>2120 (R2120)</i>
	SSC by employees, payroll basis	<i>2210 (R2210)</i>
	SSC by employees, income tax basis	<i>2220 (R2220)</i>
	SSC by self-employed, payroll basis	<i>2310 (R2310)</i>
	SSC by self-employed, income tax basis	<i>2320 (R2320)</i>
Payroll (3000), 1 position	SSC unallocable	<i>2400 (R2400)</i>
	Taxes on payroll and workforce (non-SSC)	<i>3000 (R3000)</i>
Property & wealth (4000), 10 positions	Recurrent taxes on immovable property, households	<i>4110 (R4110)</i>
	Recurrent taxes on immovable property, other	<i>4120 (R4120)</i>
	Non-recurrent taxes on	<i>4210 (R4210)</i>

	financial net wealth	
	Non-recurrent taxes on non-financial net wealth	4220 (R4220)
	Recurrent taxes on net wealth of individuals	4310 (R4310)
	Recurrent taxes on net wealth of corporations	4320 (R4320)
	Estate, inheritance, and gift taxes	4400 (R4400)
	Taxes on financial transactions	4510 (R4510)
	Taxes on capital transactions	4520 (R4520)
Consumption (5000), 4 positions	Other recurrent taxes on property	4600 (R4600)
	General taxes on goods and services (VAT)	5110 (R5110)
	Taxes on specific goods and services (excises)	5120 (R5120)
	Unallocable between general and specific taxes	5130 (R5130)
	Taxes on use of goods / permission to perform activities	5300 (R5300)
Other (6000), 1 position	Other taxes not elsewhere classified	6000 (R6000)

Notes: $K = 28$ positions. Revenue shares sum to one for each country-year and define each country's vector in the cosine-similarity network S . Tax diversification is the normalised Shannon entropy of country i 's row in S , $H = -\sum_j s_{ij} \ln(s_{ij})/\ln(28)$, where the 28 in the denominator is the number of peer countries (29 economies less the diagonal self-loop), not the 28 tax positions; the positions enter only by defining the revenue-share vectors from which the similarities are computed. PIT separated from CIT (R1110–1120 vs R1210–1220) because of systematically different distributional properties. R4600 (Other recurrent taxes on property) is the tenth property position. Source: OECD Revenue Statistics and interpretative guide (OECD, 2025a, 2025b).

Appendix B. Country Structural Positions (sorted by EIG descending)

Table B1. Country eigenvector centrality and network position (1995–2022 medians)

Rank	Country	EIG (norm.)	Region	Backbone links	Top peer	Top sim.
1	Finland	1.000	Northern	9	Switzerland	0.891
2	Switzerland	0.997	Western	8	Finland	0.891

3	France	0.945	Western	6	Finland	0.866
4	Croatia	0.940	Eastern	7	Switzerland	0.819
5	Italy	0.907	Southern	5	France	0.821
6	Austria	0.907	Western	5	Croatia	0.795
7	Spain	0.900	Southern	4	France	0.807
8	Slovenia	0.870	Eastern	4	Croatia	0.762
9	Denmark	0.855	Northern	3	Switzerland	0.756
10	Netherlands	0.850	Western	3	Finland	0.744
11	Ireland	0.834	Western	3	Switzerland	0.768
12	Slovak Republic	0.832	Eastern	3	Finland	0.728
13	Belgium	0.782	Western	3	France	0.706
14	Bulgaria	0.764	Eastern	2	Switzerland	0.696
15	Iceland	0.703	Northern	2	Finland	0.599
16	Greece	0.682	Southern	2	Switzerland	0.661
17	Latvia	0.659	Eastern	3	Lithuania	0.867
18	Poland	0.653	Eastern	2	Croatia	0.622
19	Norway	0.649	Northern	2	Croatia	0.565
20	Estonia	0.604	Eastern	2	Slovenia	0.677
21	Czechia	0.502	Eastern	2	Spain	0.468
22	Portugal	0.500	Southern	2	Spain	0.421
23	Malta	0.489	Southern	2	Denmark	0.419
24	Lithuania	0.470	Eastern	2	Latvia	0.867
25	Germany	0.460	Western	2	Croatia	0.373
26	Romania	0.388	Eastern	2	Switzerland	0.341
27	Hungary	0.349	Eastern	2	Switzerland	0.307
28	Sweden	0.338	Northern	2	Italy	0.325
29	Luxembourg	0.263	Western	2	Belgium	0.284

Notes: EIG = normalised eigenvector centrality from the full 29×29 cosine-similarity matrix (full-sample median). Backbone links = undirected edges retained under row-wise 95th-percentile pruning. Sorted strictly descending by EIG.

Appendix C. CIPS/CADF Panel Unit-Root Tests

Table C1. CIPS/CADF p-values (lags 0–3, intercept and trend)

Variable	Intercept (lags 0–3)				Trend (lags 0–3)				Δ (lag 0)
	L0	L1	L2	L3	L0	L1	L2	L3	p
Gini (disp.)	0.239	0.014*	0.297	0.239	0.991	0.514	0.920	0.701	0.000** *
<i>Tax diversification</i>	0.000** *	0.013*	0.003**	0.553	0.000** *	0.002**	0.000** *	0.433	0.000** *
<i>Fiscal centrality</i>	0.010*	0.214			0.001** *	0.398			0.000** *
<i>VAT gap (lagged)</i>	0.934	0.836			0.836				0.000** *
<i>Log GDP pc</i>	0.000** *	0.000** *	0.000** *	0.074 †	0.030*	0.030*	0.240	0.240	0.000** *
<i>Unemployment</i>	0.665	0.011*	0.028*	0.757	0.582	0.000** *	0.000** *	0.000** *	0.000** *
Gov. effectiveness	0.255	0.342	0.987	0.831	0.189	0.174	0.981	0.655	0.000** *

Notes: H_0 : unit root. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$, † $p < 0.10$ marginal.

Appendix D. Mechanism Check Results

Table D1. Nonlinear diversification (M5) and institutional moderation (M7)

	M5: nonlinear diversification	M7: institutional moderation
L.Gini (ϕ)	0.887*** (0.031)	0.907*** (0.027)
<i>c_sel</i> (diversification, centred)	−0.0140*** (0.0019)	−0.0173*** (0.0020)
<i>c_sel</i> ² (quadratic term)	+0.000137*** (0.000037)	—
<i>Fiscal-network centrality</i>	0.0123*** (0.0017)	0.0104*** (0.0013)
<i>c_sel</i> × <i>c_ge</i> (interaction)	—	+0.000530*** (0.000134)
W·Gini (ψ)	0.065† (0.036)	0.053 (0.046)
W·VAT gap	0.0060** (0.0022)	0.0022 (0.0026)
Test of mechanism term (p)	$c_sel^2 = 0$: p = 0.0002	$c_sel \times c_ge = 0$: p = 0.0001
Hansen J (p)	0.366	0.626
N	754	754

Notes: $c_sel = sel - \text{mean}(sel)$; $c_ge = ge - \text{mean}(ge)$. M5: positive quadratic implies diminishing returns to diversification at high levels; the marginal effect remains negative throughout the observed data range. M7: positive

interaction implies that institutional quality and tax diversification are partial substitutes, countries with weaker governance derive greater inequality reduction from diversification. *** $p < 0.01$, ** $p < 0.05$, † $p < 0.10$.

Appendix E. Weight-Matrix Construction

All weight matrices are built from the pairwise cosine similarity of the 28-position revenue-share vectors, pruned and row-normalised before estimation.

Matrix	Basis	Pruning	Row-normalised	Used for
W1 (W _{sim})	cosine similarity	95th percentile	yes	baseline
W2 to W5	cosine similarity	90th / 85th / 80th / 75th percentile	yes	threshold robustness
W1R to W5R	1 – cosine similarity (dissimilarity)	5th / 10th / 15th / 20th / 25th percentile	yes	dissimilarity robustness
WEARLY	first-half median similarity	95th percentile	yes	subperiod robustness
WLATE	second-half median similarity	95th percentile	yes	subperiod robustness

Appendix F. Network Construction and the Similarity-Space Diversification Measure

F.1. Two objects: the information network S and the topology matrix W

The analysis uses the fiscal-similarity network for two analytically distinct purposes. The information object is the network S itself: for each year, the cosine similarity between every pair of countries' 28-position revenue-share vectors gives a dense 29-by-29 similarity matrix, the diagonal self-loop is removed, and the node statistics used as regressors are computed on this full matrix. The topology object is the matrix W: the same annual similarity matrices are pruned at a similarity threshold and row-normalised, and the element-wise full-sample median is used as the fixed interaction matrix that defines the spatial lag and contextual terms in the defactored IV estimator. Both derive from the same cosine-similarity construction, but they serve different roles. S carries the information used to characterise each country's position; W carries only the topology used to define fiscal neighbours in estimation. Computing a node statistic on the pruned W rather than on the full S would be both technically incorrect and economically uninterpretable, because the pruned matrix retains only one or two links per country and discards the structure that the node statistics summarise.

F.2. The tax-diversification measure

Tax diversification is the normalised Shannon entropy of country i's row in the annual similarity network S, $H = -(1/\ln 28) \sum_j s_{ij} \ln s_{ij}$, where the s_{ij} are the row-normalised similarity weights from country i to its 28 European peers and the $\ln(28)$ normalisation refers to the 28 peers (29 economies less the diagonal self-loop), not to the 28 tax positions. The 28 positions enter only by defining the revenue-share vectors from which the pairwise similarities are computed. The index is high when a country resembles many peers broadly and evenly and low when its resemblance concentrates on a small cluster. This similarity-space entropy is the variable used in the regressions: rescaled to [1.00001, 100] it reproduces the Table 1 mean of 80.04 exactly. The within-country entropy of the 28 revenue shares directly is a different object, rescaling to a mean of 59.2 and correlating with the variable used at 0.46 (Spearman 0.31); it is not the variable

estimated. Because a country's similarity position is generated by its own revenue composition, the two measures move together in part, and diversification remains a domestically shaped margin even though it is measured in fiscal-similarity space. As a robustness check (Section 6, Table 6), the preferred specification re-estimated with the direct 28-share entropy in place of the similarity-row entropy returns a coefficient of -0.0033 ($p = 0.015$), negative and significant, so the inequality-reducing association is not an artefact of the network-based construction.

F.3. Stability of the network statistics

The node statistics are stable across the 28 annual networks, which justifies treating the interaction matrix as a fixed structural object. Eigenvector centrality has a median within-country coefficient of variation of 1.8%, a year-on-year rank correlation of 0.96, and a rank correlation of 0.81 across all year-pairs. The similarity-row entropy varies by under 0.1% within countries, with a year-on-year correlation of 0.95. Section 6 re-estimates the model on subperiod networks built separately from the first and second halves of the sample as a further safeguard against secular structural change. The country eigenvector-centrality values reported in Appendix B are a post-hoc summary computed from the data and are not a direct read-off of the dense annual adjacency matrices.