International Tax and Public Finance
https://doi.org/10.1007/s10797-026-09976-2

Fiscal policy divergence

Piera Bello[1] · **Sergio Galletta**[2] · **Costanza Marconi**[1]

Received: 9 October 2025 / Accepted: 6 May 2026

Abstract
We propose a novel measure of fiscal policy divergence based on public budget data. First, we quantify policy similarity across municipalities based on the cosine similarity of their budget allocations, showing strong correlations with geographic proximity and differences in socio-demographic characteristics. Next, we predict budget similarity out-of-sample using a flexible machine-learning model trained on pairwise differences in local characteristics, defining fiscal divergence as the municipality-year average absolute prediction error. In an empirical application to 8000 Italian municipalities (2000–2015), we show that electoral-cycle patterns are concentrated on the expenditure side of the budget: expenditure divergence is lower in pre-electoral years than in election years, whereas revenue divergence follows a different and less stable pattern.

Keywords Fiscal policy · Policymaking · Italy · High-dimensional data

JEL Classification H72 · H20 · H30

For helpful comments and suggestions, we thank seminar participants at the University of Pavia, LISER, Bocconi, ETH Zurich (Economics + Data Science series). We also thank participants at the 2023 SIEP Conference (Verona), 5th Swiss Workshop on Local Public Finance and Regional Economics (Lugano), 2nd Workshop on The Political Economy of Municipal Fiscal Policy (Bergamo), 2024 ZEW Public Finance Conference (Mannheim), and 2024 IIPF Annual Congress (Prague). P. Bello and C. Marconi acknowledge that this project is within the GRINS – Growing Resilient, INclusive and Sustainable Extended Partnership – project, funded by the European Union –NextGenerationEU, under the National Recovery and Resilience Plan (NRRP).

✉ Sergio Galletta
sergio.galletta@uniroma1.it

Piera Bello
piera.bello@unibg.it

Costanza Marconi
costanza.marconi@unibg.it

[1] Università di Bergamo, Bergamo, Italy

[2] Sapienza University of Rome, Rome, Italy

Published online: 21 June 2026

1 Introduction

Fiscal policy is a crucial tool for local governments, shaping their capacity to deliver essential services and promote economic growth within their jurisdictions. Ideally, the allocation of expenditure and generation of revenue should be structured to address local needs while prioritizing public welfare (Oates, 1972). Indeed, if local preferences and demographics significantly influence the optimal policy choices (Agrawal et al., 2022), fiscal policy should mirror the specific characteristics of the local community. However, evaluating how well a fiscal policy is consistent with the local population's needs poses two related empirical challenges. First, even if local preferences and demographics shape optimal policy choices, there is typically no clear benchmark against which a jurisdiction's fiscal policy can be assessed. Without a standard for comparison, it is difficult to determine whether a given budget reflects local needs or represents an unexpected departure from them. Second, local fiscal policy is inherently multidimensional. Local governments simultaneously choose among multiple taxes and spending categories, offering residents a package deal rather than a single policy instrument. As emphasized by Agrawal et al. (2022), focusing on individual taxes or expenditures in isolation may therefore provide an incomplete characterization of local policy choice, particularly when jurisdictions can substitute across fiscal instruments. This issue is especially salient in institutional settings where multiple revenue sources and spending categories coexist, and where adjustments along one margin can be offset by changes elsewhere in the budget (see, e.g., Agrawal et al. 2025).[1]

We propose a novel approach that addresses both challenges simultaneously. We construct a comparative benchmark from socio-demographic similarity across jurisdictions and characterize fiscal policy using the full composition of the municipal budget. By representing budgets as high-dimensional objects, we measure policy similarity and divergence across the joint fiscal package without imposing ex ante restrictions on which margins matter or arbitrary weights on specific instruments.

The standard public finance literature emphasizes neighboring jurisdictions' policies as accessible benchmarks for voters and politicians (Besley & Case, 1995). Our intuition is similar, but we replace geographic proximity as the sole benchmark with observable similarity in socio-demographic characteristics. We then ask whether pairs of municipalities that look similar in these characteristics also make similar fiscal choices. Departures from the predicted level of similarity provide a proxy for unexpected local fiscal behavior. We remain agnostic about the source or normative meaning of such divergence: it may capture innovative policymaking (Bernecker et al., 2021), but it may also reflect inefficiencies, mismanagement, or corruption (Bandiera et al., 2009). Our contribution is to provide a consistent way to detect and quantify these departures, leaving their interpretation to the context of specific applications.

[1] Agrawal et al. (2025) use a composite index to summarize the joint fiscal stance of local governments engaged in intermunicipal cooperation, as in France, where municipalities levy multiple capital taxes, making the analysis of any single tax potentially misleading.

To demonstrate our approach, we draw on detailed data on public budgets and socio-demographic characteristics for roughly 8000 Italian municipalities between 2000 and 2015. Municipal budgets provide a reliable proxy for government policy, as they reflect choices over spending, revenue, and resource allocation. Our first contribution is to construct a comprehensive measure of budget similarity across municipalities by analyzing the entire structure of local budgets with tools from applied machine learning. We represent each municipality's budget as a high-dimensional vector–spanning approximately 4400 categories–and compute similarity as the cosine of the angle between these vectors. A higher cosine similarity indicates that two municipalities allocate resources in more comparable ways, capturing the closeness of their budgetary choices. Importantly, we show that municipalities that are geographically proximate, share similar population structures, or are led by mayors from the same political party exhibit markedly higher budget similarity. Among these factors, geographic distance and population age composition stand out as the strongest predictors. By contrast, individual attributes of mayors–such as age or gender–have little explanatory power, suggesting that local fiscal choices are shaped more by structural and political contexts than by personal characteristics of officeholders.

We next assess how actual budgets diverge from the similarity predicted by municipal characteristics. The purpose of this prediction step is not to forecast municipal budgets perfectly, but to construct an empirical benchmark for expected fiscal similarity. We estimate the benchmark separately by region and year, using absolute pairwise differences in a rich set of local characteristics as predictors of budget similarity. Our preferred specification relies on gradient-boosted trees (XGBoost) estimated in 5-fold cross-validation, so that every similarity prediction is evaluated strictly out-of-sample. For the full-budget measure, the average out-of-sample R^2 is 0.189. This is meaningful predictive content for a noisy fiscal outcome, while leaving substantial unexplained variation. We use the absolute out-of-sample prediction error as the measure of fiscal divergence: larger errors indicate budget allocations that depart more strongly from demography-based expectations.

To provide a practical example of how our measure can be applied, we study the role of electoral incentives in shaping local fiscal choices through the political budget cycle. Specifically, we test whether fiscal divergence from demography-based expectations varies systematically over the electoral cycle, exploiting the staggered timing of municipal elections across Italian municipalities (Ferraresi, 2020; Repetto, 2018). The results differ across fiscal blocks. Expenditure divergence is lower before elections, while revenue divergence does not display the same pattern. This is consistent with models in which politicians adjust the composition of visible public spending to signal responsiveness to voter priorities (Drazen & Eslava, 2010), and with evidence that local political budget cycles often operate through the spending side rather than through all fiscal instruments equally.

More broadly, this application illustrates how our measure of fiscal divergence can be used in a variety of settings to study determinants and consequences of local policy choices. Beyond electoral incentives, the framework can be employed to analyze policy innovation and diffusion across jurisdictions, the role of political alignment or decentralization reforms, the fiscal consequences of the rise of populist parties at the

local level, or the relationship between unexpected fiscal choices and outcomes such as corruption, fiscal stress, or citizen satisfaction.

We contribute to the body of literature on fiscal policy and its drivers. The determinants of fiscal policy include the characteristics of the mayor and the ruling party (Brollo & Troiano, 2016; Ferreira & Gyourko, 2014); media exposure, such as the influence of conservative or progressive media channels on voters' preferences, which can be reflected in diverging budget allocation choices (Ash & Galletta, 2023); the degree of decentralization or fiscal autonomy enjoyed by local policymakers (Daniele & Giommoni, 2025; Galletta, 2020; Grembi et al., 2016); electoral rules and term limits (Gagliarducci & Nannicini, 2013); and monetary incentives (Ferraz & Finan, 2009). Moreover, local economic shocks, such as trade-induced changes in the local tax base (Feler & Senses, 2017) or broader economic disruptions (Janeba et al., 2025), as well as fiscal shocks, including unexpected changes in intergovernmental transfers (Helm & Stuhler, 2024; Martínez, 2025), are shown to significantly affect local governments' fiscal decisions. In addition, local governments' actions are partly driven by strategic considerations arising from fiscal competition with neighboring jurisdictions (Agrawal et al., 2022; Janeba & Osterloh, 2013). Unlike previous studies that typically focus on a single component of local policymaking, our method accounts for the overall budget composition, thereby improving the analysis of fiscal decisions.

Finally, we provide additional evidence supporting the theory of a political business cycle (Nordhaus, 1975). Most evidence indicates that under certain conditions, politicians manipulate budgets such that, closer to elections, public expenditure increases while taxes might decrease. This often results in a negative impact on public debt and deficits (Akhmedov & Zhuravskaya, 2004; Alesina & Paradisi, 2017; Ferraresi, 2020; Repetto, 2018). Other models suggest that politicians adjust the composition of spending, not the overall budget, to signal their priorities to voters (Drazen & Eslava, 2010).

The remainder of the paper is organized as follows. Section 2 provides the institutional background and presents the data. Section 3 details the methodology for computing budget similarity and constructing pairwise local-characteristic differences, and tests the validity of our metric. Section 4 introduces and estimates our measure of policy divergence. Section 5 discusses the results of the empirical application concerning the political budget cycle. Finally, Section 6 offers concluding remarks and summarizes the key findings.

2 Institutional background and data

2.1 Institutional background

Italy's sub-national government features a three-tier system: 20 regions (regioni), 110 provinces (province), and approximately 8,000 municipalities (comuni). This study focuses on the municipal level, the lowest administrative unit. Each municipality (comune) has a mayor (sindaco), an executive committee (giunta) appointed by

the mayor, and an elected city council (consiglio comunale) that approves the annual budget proposed by the mayor.

Municipalities in Italy are responsible for several public services, such as waste disposal, local transportation, social services, childcare and primary schooling, urban road maintenance and cleaning, water and sewer services, environmental monitoring and protection, planning and zoning. Municipal revenues consist of various sources, including tax revenues derived from income taxes, real estate taxes, and taxes related to services like waste management. Additionally, transfer revenues originate from the national or regional governments, as well as from the European Union (EU). On the expenditure side, current expenses pertain to the municipality's day-to-day operational costs, such as salaries and utilities. Capital expenditures, on the other hand, are investments allocated to projects that typically extend beyond a single budget year and are primarily associated with infrastructure development, such as the construction of roads and schools.

The degree of municipal discretion varies considerably across these budget items. On the revenue side, Italian municipalities have direct control over the rates of the local property tax (ICI/IMU), the surtax on personal income (*addizionale comunale all'IRPEF*), and a number of fees and tariffs (waste collection, building permits and related sanctions, school services), within bands set by national legislation. Property-tax bases, building-permit revenues, and own fees are therefore the most clearly discretionary revenue components. By contrast, intergovernmental transfers from the State, the regions, the EU, and other public bodies–including the historically important *Fondo Ordinario* and the more recent equalization funds–are determined by national or regional formulas tied to demographic and fiscal indicators, leaving municipalities with little control over their level. On the expenditure side, capital spending on roads, traffic, and land-use planning, current spending on general administration, and a portion of operating service contracts (waste management, local transport) are largely under municipal control, especially through procurement and staffing decisions. Other categories–most notably earmarked social and education spending (childcare, residential facilities for the elderly, school assistance, transport, canteens)–are partly delegated from higher levels of government and financed by earmarked transfers, so the municipality typically chooses how, but not whether, to provide the service. We use this distinction in two ways. First, the baseline similarity index uses the entire fiscal package precisely because municipalities can substitute across margins and the resulting policy mix is a relevant economic object even where each individual item is not fully discretionary. Second, in the Appendix we report a robustness exercise (Figure 10 and Section Appendix A.2) that re-computes the similarity index after dropping the items most plausibly outside the influence of the local government.

The current framework for municipal elections in Italy applies the direct election of mayors and the plurality rule, with variations based on city size. In municipalities with populations below 15,000, elections use a single ballot and plurality rule. In cities with a population exceeding 15,000, a dual ballot system is employed. Since

1993, mayors in Italy have been subject to a two-term limit.[2] In 2000, the length of the mayoral term was prolonged from four to five years. The number of city councilors varies based on the size of the municipality.

2.2 Data

We compiled a novel dataset from several sources. First, we gathered municipal budget data for the 2000–2015 period from the Ministry of Economy and Finance for all 8000 Italian municipalities. Our dataset provides information on the complete composition of local budgets. In particular, we use data on both revenue and expenditure.[3] For expenditures, each variable specifies whether they are commitments, payments in current account competence, or payments in residual accounts. Similarly, for revenues, we account for assessments, collections in current account competence, and collections in residual accounts. While the budget composition changes each year slightly, it remains consistent across municipalities for any given year. Notably, there were significant changes in 2008, which we will detail in our analysis. As a reference, in 2008, we have 4413 components in total. Of these, 294 come from revenues, 1977 from current expenditure, and 2142 from capital expenditure. The largest components are building permits and related sanctions, the ordinary grant fund, and local taxes on the revenue side, while expenditure is mainly concentrated in general administration and personnel costs for current spending and in roads, traffic, and land-use planning for capital spending. We provide details on the budget structure and largest components in Appendix A.

The second main source is the 2011 Italian decennial census, which we use as a cross-sectional snapshot of the demographic and economic composition of each municipality. We include variables related to housing and living conditions, mobility and transportation, economic indicators, environmental and urban quality, and demographic and migration dynamics. The complete list of these variables is presented in Appendix Table 2. All socio-demographic variables listed in Table 2 are constructed from the 2011 Census and are treated as time-invariant in the empirical analysis. In principle, a subset of variables (housing characteristics, income brackets, and migration indicators) could have been built in a time-varying fashion using yearly registries, but for transparency and consistency we keep the entire socio-demographic block fixed at its 2011 value.

[2] Since 2014 a reform (Law April 2014 no.56) allows mayors in municipalities with less than 3000 inhabitants to re-run for a third term, whereas mayors in cities with a number of residents above the cut-off still face a two-term limit.

[3] Italian local budgets are divided into sections (*quadri*). Our data is sourced from *quadro 2* (revenues), *quadro 4* (current expenditures), and *quadro 5* (capital expenditures). We do not include items related to municipal borrowing and debt repayment, which appear in *quadro 3* and are further detailed in *quadri* 8, 10, and 12. The reason is that debt and borrowing items in Italian municipal budgets have limited cross-sectional variation that genuinely reflects local policy: under Article 119 of the Constitution, municipalities can only borrow to finance investment, not current spending, and the Internal Stability Pact (*Patto di Stabilità Interno*) imposed binding fiscal targets on the bulk of municipalities throughout our sample (1999–2015). Including these items in the similarity index would therefore largely capture compliance with national fiscal rules rather than discretionary local choices.

Finally, we incorporated information about local politics from the "Anagrafe degli Amministratori Locali e Regionali" database of the Ministry of Interior. This dataset includes details on the gender, educational attainment, and party affiliation of the mayor, as well as the year of elections.

3 Similarity in fiscal policy and local characteristics

3.1 The similarity index

In this section, we introduce our measure of fiscal-policy similarity and the local characteristics used to explain it. The outcome of interest is a metric that shows how similar the policymaking of a pair of municipalities is. Because local policy is not captured only by the size of the budget or by allocations in broad macro categories, we represent each municipal budget as a fine-grained vector of revenue and expenditure shares. Drawing from recent applications of high-dimensional data methods, we use cosine similarity to compare these budget vectors. Local socio-demographic and political characteristics enter the analysis separately as standardized pairwise differences.

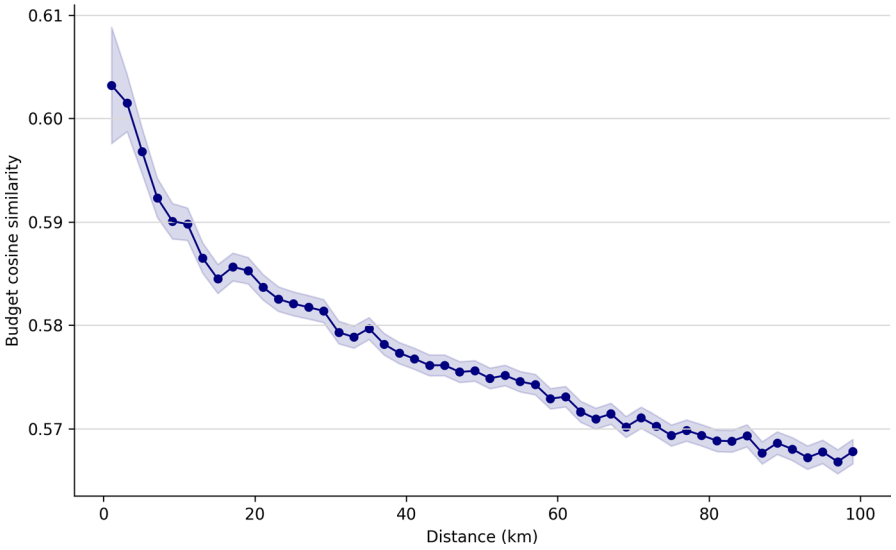

Fig. 1 Binscatter budget similarity vs geographical distance. The binscatter plot illustrates the relationship between the cosine similarity of municipal budgets and the geographical distance between municipalities. Each point represents an aggregated average of budget similarity and distance (in km) across municipality pairs falling in a given distance bin; the shaded area is the 95% confidence interval around the local mean

This similarity metric measures the cosine of the angle between two non-zero vectors in a multidimensional space. It assesses how similar the two vectors are by focusing on their direction rather than their magnitude. The value ranges from − 1

to 1, where 1 indicates identical direction, 0 indicates orthogonality (no similarity), and − 1 indicates completely opposite direction. However, when all vectors have elements that are non-negative, as in our case for the budget, this measure of similarity is bounded in the interval [0,1].

$$\cos(\theta) = \frac{\mathrm{A} \cdot \mathrm{B}}{\|\mathrm{A}\| \, \|\mathrm{B}\|} = \frac{\sum_{i=1}^{n} a_i \cdot b_i}{\sqrt{\sum_{i=1}^{n} (a_i)^2} \sqrt{\sum_{i=1}^{n} (b_i)^2}} \tag{1}$$

Formally, we compute the dot product between vectors A and B and divide it by the product of their magnitudes. The dot product is the sum of the products of corresponding elements (a_i and b_i) of the two vectors. The magnitudes are calculated as the square root of the sum of the squares of the elements of each vector as reported in Equation (1).[4]

For our purposes, the vectors of interest for assessing similarity in policymaking are derived from the budget accounts, where each item is expressed as a proportion of the relevant total revenue or expenditure, depending on the item's category. This means that we have one fiscal vector for each municipality-year. The local-characteristic variables are instead used to form pairwise covariate differences between municipalities. All socio-demographic variables in Table 2 come from the 2011 decennial census and are treated as time-invariant; pairwise covariate differences therefore vary across pairs but not across years. Vectors (2) and (3) illustrate the budget components used to compute fiscal similarity and the local characteristics used to construct pairwise predictors.

$$\begin{pmatrix} \text{Budget Vector} \\ \text{Share Taxes (Rev.)} \\ \text{Share Prop. Tax (Rev.)} \\ \text{Share Prop. Tax First Res. (Rev.)} \\ \text{Share Prop. Tax Oth. Buil. (Rev.)} \\ \vdots \\ \vdots \\ \text{Share Acqui. of Real Estate (Exp)} \\ \text{Share Acqui. Machinery (Exp)} \\ \text{Share Ext. Prof. Assig. (Exp)} \\ \text{Share Capital Contr. (Exp)} \end{pmatrix} \tag{2}$$

[4] It is worth highlighting that the cosine similarity formula is closely related to the Pearson correlation coefficient when applied to centered vectors (variables).

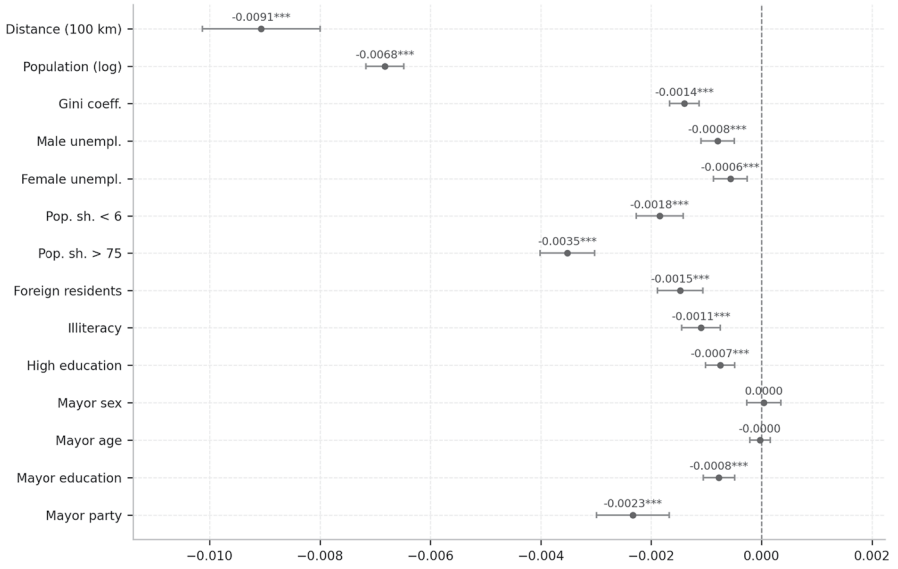

Fig. 2 Validation policy similarity. This graph shows the coefficients and confidence intervals of various municipal characteristics differences between a pair of municipalities (standardized) on their fiscal policy similarity. The stars denote significance levels: ***p < 0.01, **p < 0.05, *p < 0.10. The x-axis shows the magnitude of the coefficients, while the y-axis lists the variables. The error bars represent the 95% confidence intervals for each coefficient

$$\begin{pmatrix} \text{Local Characteristics} \\ \text{Population} \\ \text{Share of Employed} \\ \text{Share of highly educated} \\ \text{Share of foreigners} \\ \vdots \\ \vdots \\ \text{Population Density} \\ \text{Literacy Rate} \\ \text{School Dropout} \\ \text{Divorce rate} \end{pmatrix} \tag{3}$$

It is important to highlight that by using vector representation and cosine similarity for the budget measure, we abstract away from potential differences in the relative importance of each component in policy formulation.[5]

[5] In this approach, each component is treated with equal weight, implying that all elements contribute equally to the overall similarity measure. This approach might introduce "measurement error" in the similarity estimates, as we could underweight highly relevant components or overweight less relevant ones. We abstain from introducing differential weights as this could be seen as arbitrary, unless there is a clear rationale for assigning specific importance to particular components.

3.2 Generating budget similarity and pairwise covariates

We construct one fiscal-policy similarity measure and a corresponding set of pairwise local-characteristic differences. The fiscal-policy measure captures similarities in policymaking by analyzing the detailed structure of municipal budgets, reflecting local governments' fiscal decisions and priorities. The local-characteristic variables summarize how each municipality pair differs in population structure, socioeconomic conditions, geography, and mayoral attributes.

To generate similarities in policymaking, we analyze the complete structure of the budgets.[6] The number of components differs across years; Table 1 offers a detailed overview of the final number of budget variables employed in each year. Overall, we have more information on the expenditure side rather than the revenue side. We can also see a marked increase in the number of expenditure features in the year 2008. This variation in budgetary components is important to keep in mind, as it implies that the generated similarity measure primarily offers a cross-sectional interpretation of policy similarity. We restrict the construction of similarity measures to pairs of municipalities within the same region and the same year, both for computational and substantive reasons. Computationally, computing all pairwise similarities for the roughly 8000 Italian municipalities across 16 years would generate hundreds of millions of pair-year observations and would make subsequent regressions with municipality fixed effects very demanding; the within-region restriction reduces the data by an order of magnitude. Substantively, Italian regions differ considerably in their institutional frameworks and fiscal autonomy – particularly the five *regioni a statuto speciale* – and these differences would introduce systematic variation in budget structures that is unrelated to local policy choices. Restricting comparisons to within-region pairs ensures that we compare municipalities operating under the same institutional environment, and is internally consistent with the rest of our analysis: the divergence model in Equation (5) is estimated separately by region and year, and the electoral-cycle analysis in Section 5 further absorbs region-level heterogeneity through municipality fixed effects.

The increase in the number of budget components after 2008 is relevant for interpreting comparisons over time. For this reason, our validation and prediction exercises focus on within-year and within-region variation rather than on unconditional changes in the average level of budget similarity.

To construct the local-characteristic predictors, we account for variables related to housing and living conditions, mobility and transportation, economic indicators, environmental and urban quality, and demographic and migration dynamics. We use 60 variables relevant to describing a municipality's population structure (see Appendix Table 2). For each municipality pair, the prediction model uses standardized absolute differences in these characteristics. Variables are standardized within

[6] We exclude only those variables that are missing for all municipalities in a given year. Such uniform missingness arises primarily from the institutional structure of the budget classification, which can change from one year to another for political or administrative reasons: tax classifications evolve over time, leading some denominations to become obsolete, and certain line items are only relevant for specific jurisdictions or are activated in particular periods. As a result, a variable can be missing for all municipalities in a given year when that specific budget item is not used or recorded.

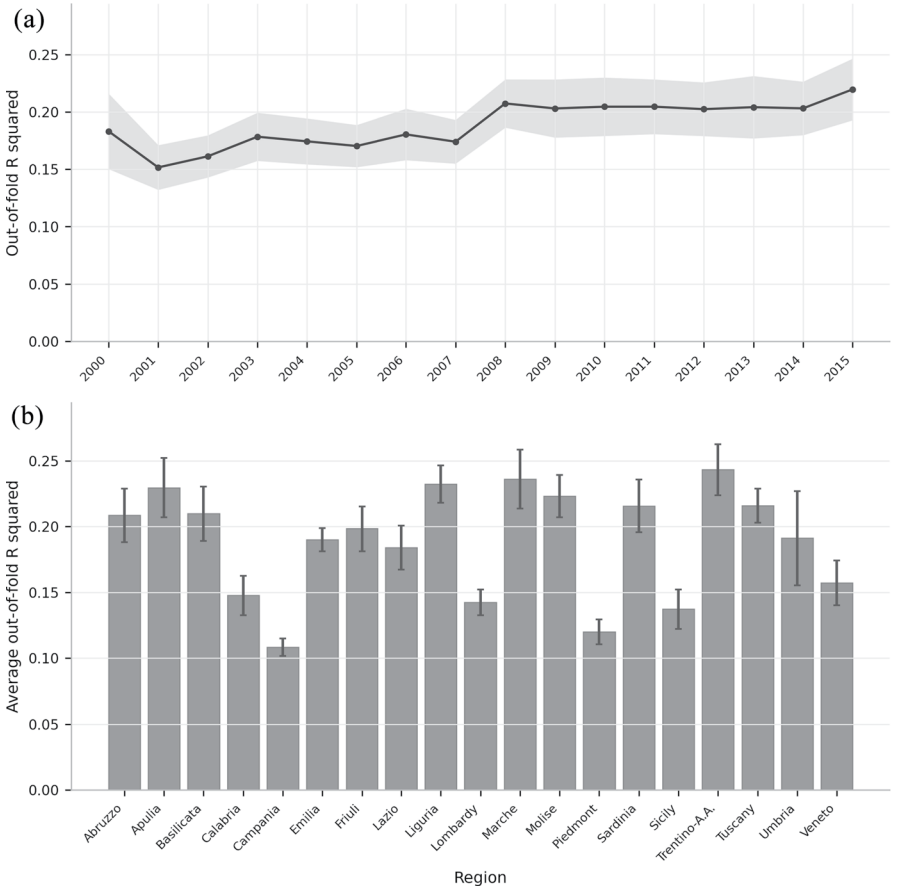

Fig. 3 Out-of-sample prediction performance - time and regional variation. *Notes:* This figure reports the temporal and regional variation in out-of-sample prediction performance for the full-budget measure. Panel (a) shows the average R^2 by year across regions, while panel (b) shows the average R^2 by region across years. Shaded areas and error bars denote 95% confidence intervals. Appendix figures report the corresponding plots for revenue and expenditure

region, which is the level at which we ultimately estimate the prediction model used to calculate the divergence measure.

3.3 A validation exercise

In this section, we provide evidence supporting the validity of our measure of policymaking similarity by analyzing how it relates to differences in socio-demographic, geographic, and political characteristics known to be important determinants of policymaking. By examining these relationships, we aim to demonstrate that municipalities with more similar profiles tend to have more similar budget structures, thereby validating our budget cosine index as an effective measure of policymaking similarity.

We begin our analysis by examining the relationship between the cosine similarity index and geographical distance, a critical factor discussed in public finance

literature.[7] Figure 1 presents a binscatter plot that illustrates a negative correlation between budget similarity and the distance between pairs of municipalities. This finding supports the hypothesis that geographical proximity influences budgetary decisions. Municipalities that are farther apart tend to exhibit greater differences in other characteristics, as evidenced by lower cosine similarities in their budgets.

We complement the previous evidence by estimating the following regression model:

$$\text{Policy Similarity Index}_{ij,t} = \alpha + \beta_m X^m_{ij,t} + \mu_i + \delta_j + \epsilon_{ij} \tag{4}$$

where $\text{Policy Similarity Index}_{ij,t}$ is the cosine similarity index between municipalities i and j at time t within the same region, $X^m_{ij,t}$ is the difference in the value

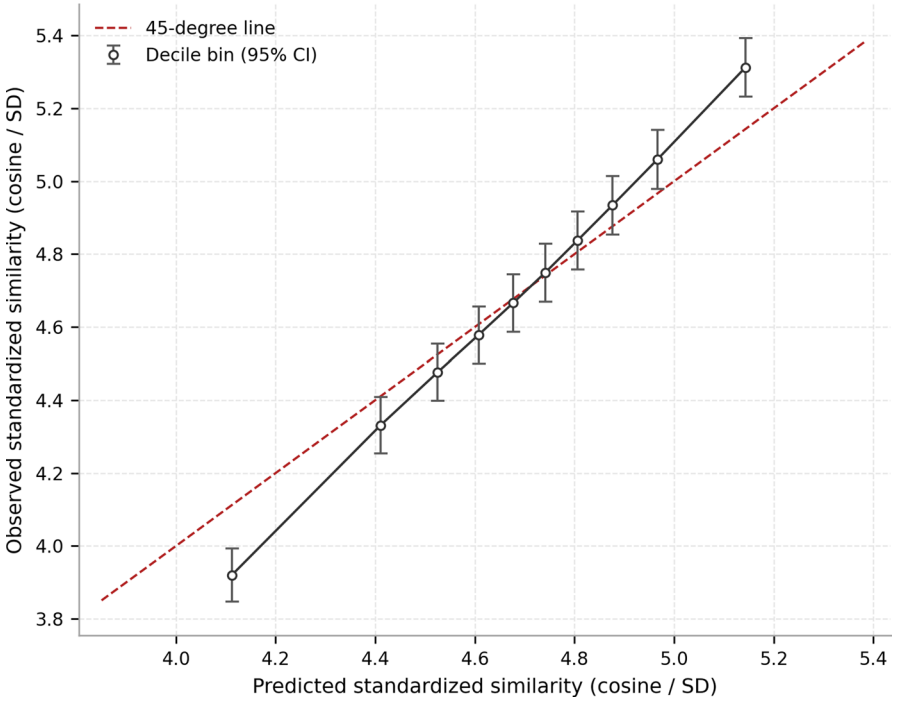

Fig. 4 Binned actual-versus-predicted budget similarity. Note: Mean observed and mean predicted standardized similarity within equally sized bins of the predicted distribution. The outcome is the standardized cosine similarity (raw cosine divided by its within-region-year standard deviation), so axes are not bounded in [0, 1]. The dashed red line is the 45-degree reference. Error bars are 95% confidence intervals across region-year cells

[7] A large body of work documents strategic interaction in local policy choices (Bordignon et al., 2003; Brueckner, 1998), tax mimicking and yardstick competition (Allers & Paul Elhorst 2005; Bosch & Solé-Ollé 2007; Buettner 2001), and expenditure spillovers (Solé-Ollé 2006), as well as spatial patterns in local efficiency and taxation (Revelli, 2001, 2007).

of our set of $m \in M$ covariates. Specifically, it represents the standardized absolute difference between each pair of municipalities of the following variables: distance (in 100 km), population (log-transformed), Gini coefficient, male and female unemployment rates, population shares under 6 and over 75 years, the share of foreign residents, illiteracy rate, share of highly educated population as well as sex, age, education, and party affiliation of the mayor. This validation specification should not be interpreted as a data-driven variable-selection exercise. Instead, we deliberately use a parsimonious, ex ante set of interpretable regressors aimed at covering broad dimensions of municipal heterogeneity that are plausibly related to differences in local policymaking. In particular, we include variables that relate to geography, municipal scale, socio-economic conditions, age structure, and demographic composition.[8] Alongside these standard predictors of local policy priorities, we include a small set of mayor characteristics in order to compare the explanatory power of municipal socio-demographic features with that of elected officials' attributes. This comparison is central to the validation exercise, as it allows us to verify whether variation in policy similarity is more strongly associated with structural local conditions or with characteristics of the political leadership.

In addition, we include municipality *i* and municipality *j* fixed effects (i.e., μ_i and δ_j). ϵ_{ij} is the error term. The municipalities' fixed effects allow us to exploit within-individual variations over different pairings. In other words, we compare how the difference in a given variable between two municipalities relates to the difference in the same variable among other municipality pairs, thereby controlling for all unobserved variables that are constant for each municipality but vary across municipalities.[9] For our exercise, we focus only on the year 2011, which is a population census year.

We present our results graphically in Figure 2. We observe that the sign of the coefficients always points in the expected direction: municipalities with greater differences in demographic and population composition, as well as mayoral characteristics, tend to exhibit lower similarities in their resource allocation and budget structure.

In particular, distance and population size emerge as the most significant predictors, emphasizing the relevance of geographic proximity and fundamental demographic characteristics. In the baseline specification, an increase of 100 km in the distance between two municipalities is associated with a decrease in budget similarity equal to about 8% of a standard deviation of the outcome distribution. The significant effects of additional population characteristics, especially age composition, inequality, and education, further highlight the role of socio-economic conditions in shaping local policy choices.

[8] Several of these dimensions–especially migration and foreign-resident composition, education, age structure, unemployment, and distance–are also among the more relevant predictors in the XGBoost prediction model introduced below. Thus, the validation exercise focuses on empirically relevant and interpretable dimensions, while the prediction model remains unrestricted and uses the broader set of available covariates.

[9] This approach is similar to the estimation of gravity models of trade in international economics, where country_i and country_j fixed effects are included in a dyadic regression. However, since our analysis does not involve time-varying characteristics, we do not include dyadic fixed effects (Anderson & Van Wincoop, 2003).

Notably, the political affiliation of mayors stands out as another important predictor of policy similarity: municipalities governed by mayors from the same party exhibit significantly more similar budget allocations, even after controlling for the differences in socio-demographic characteristics. This result suggests that shared political orientation translates into convergent fiscal choices, consistent with the expected role of partisan preferences in shaping local policymaking. Conversely, the lack of significant effects of the age and sex of mayors suggests that the individual attributes of leaders are less critical compared to their political and policy orientations.[10]

We conduct several robustness checks to assess the sensitivity of our results. First, we exclude budget items that are most likely to be mechanically driven by local characteristics, such as intergovernmental transfers and social and education expenditures (e.g., childcare services and residential care facilities for the elderly). The full-vector and restricted-vector cosine measures are very highly correlated at the pair level, with a Pearson correlation of 0.98. Re-estimating the validation specification on the restricted vector therefore yields, unsurprisingly, results that are qualitatively and quantitatively similar to the baseline (see Figure 10 in the Appendix), and the same conclusion would carry through to the divergence measure introduced in Section 4. Second, we re-estimate Equation (4) including the full set of available local characteristics. As shown in Figure 11, the coefficients of interest remain broadly stable in both direction and magnitude. We do not place structural weight on every individual coefficient in that richer specification because some migration-related controls are nearly collinear, but the main geographical and demographic gradients remain intact.

Overall, we provide supporting evidence about the validity of our budget cosine index as a measure of policymaking similarity, reflecting how comparable socio-demographic and political attributes lead to analogous budgetary decisions.

4 A measure of fiscal divergence

In this section, we construct a measure of fiscal divergence by quantifying how much municipalities deviate from the expected fiscal behavior of their socio-demographically similar peers. The intuition underlying our approach is straightforward: municipalities sharing similar characteristics are expected to make similar policy choices. Thus, substantial departures from these expected choices indicate that a municipality is pursuing unconventional or unexpected policies.

The prediction exercise is designed to build this benchmark, not to explain all variation in municipal budgets. Observable characteristics should predict the component

[10] This is interesting considering existing research that frequently identifies significant causal effects of a mayor's gender or age on various economic and social dimensions (Alesina et al., 2019; Bochenkova et al., 2023; Brollo & Troiano, 2016). The absence of significant effects of mayor age and gender should be interpreted in light of the different outcome and empirical setting considered here. While existing studies document causal effects of individual leader characteristics on specific fiscal or policy outcomes (e.g., Baskaran and Hessami, 2025; Baskaran et al., 2024), our analysis focuses on policy similarity across municipalities. In this context, geographic proximity, population characteristics, and political affiliation dominate individual attributes, suggesting that policy alignment is shaped primarily by shared environments and political preferences rather than by personal traits of officeholders.

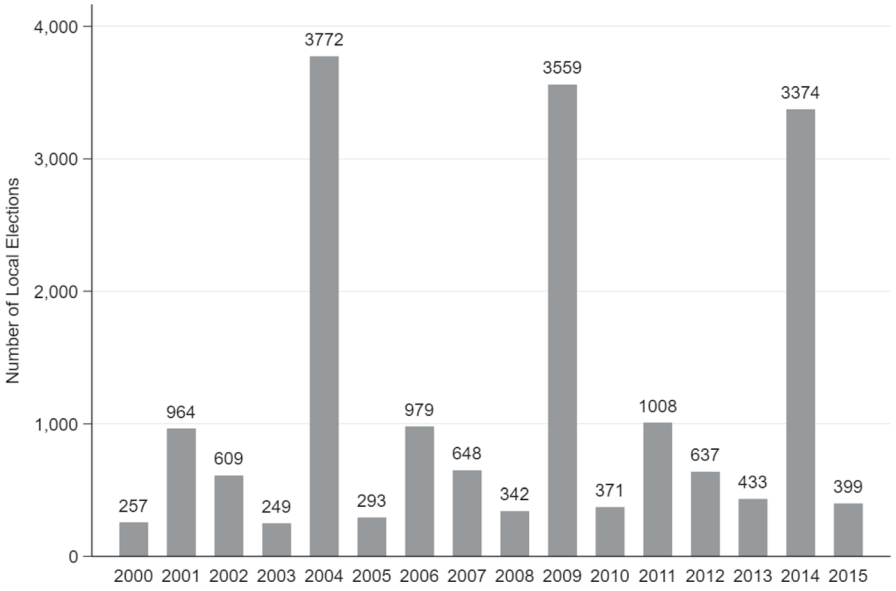

Fig. 5 Number of municipal elections by year. Note: The figure displays the number of municipal elections held each year

of fiscal similarity that reflects local needs and structural conditions; the remaining prediction error is informative precisely because municipal budgets also reflect politics, history, shocks, accounting choices, and discretion. To construct the benchmark, we move beyond the validation exercise of the previous section and use the full set of standardized pairwise covariate differences discussed in Section 3.2, which accounts for 60 variables. We estimate the following prediction rule separately for each region *r* and year *t*:

$$\text{Policy Similarity Index}_{ijrt} = \mathcal{F}_{rt}(X_{ijrt}) + \epsilon_{ijrt} \tag{5}$$

where $\text{Policy Similarity Index}_{ijrt}$ denotes the similarity in policymaking between municipalities *i* and *j* in region *r* and year *t*, X_{ijrt} is the vector of pairwise differences in local characteristics, and $\mathcal{F}_{rt}(\cdot)$ is a flexible prediction function estimated with gradient-boosted trees. We implement this prediction step using XGBoost (Chen & Guestrin, 2016), following recent applications of machine-learning methods to public finance and policy measurement (Ash & Galletta, 2024; Ash et al., 2025).

We assess prediction quality out of sample. In each region-year cell, we divide the pair-level data into five folds, iteratively hold out one fold as the test set, and train the model on the remaining four folds. For the 16 years and 19 regions in our estimation sample, this yields 1520 out-of-sample fold estimates for each budget measure.

Figure 3 summarizes the out-of-sample R^2 of the prediction model over time and across regions, and Table 3 reports complementary accuracy measures. Prediction quality varies meaningfully over time and across regions, confirming the usefulness of estimating the model separately by region and year. For the full-budget measure,

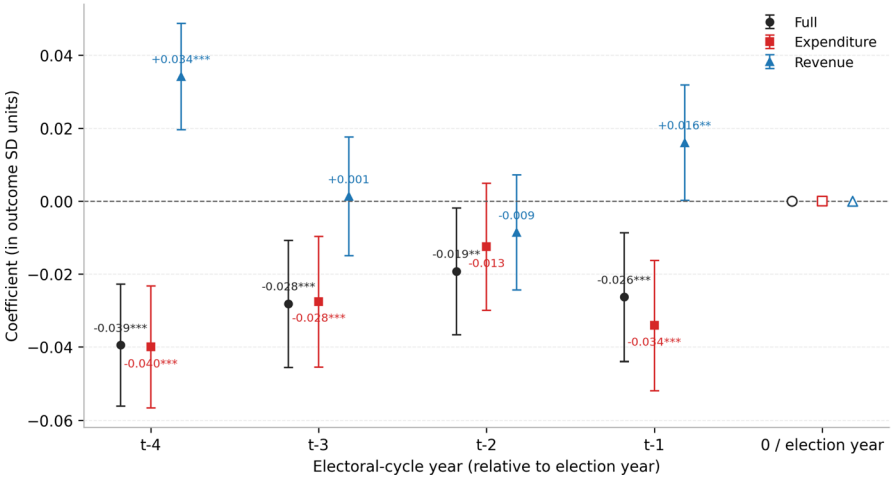

Fig. 6 Political incentives and fiscal divergence. The figure displays standardized coefficients and 95% confidence intervals from estimates of Equation (6) for the full-budget, expenditure, and revenue divergence indices. The election year ($k = 0$) is the omitted reference category and is shown as an open marker at zero on the rightmost x-axis position labeled "0 / election year". All other coefficients are interpreted relative to the election year and are expressed in outcome standard-deviation units. Markers are slightly dodged horizontally and labels are alternated above/below the points for readability; numbers next to each marker report the point estimate followed by significance stars (***p < 0.01, **p < 0.05, *p < 0.10)

the mean out-of-sample R^2 is 0.189, with an interquartile range from 0.145 to 0.227 across region-year-fold cells. The corresponding mean absolute error and root mean squared error are 0.718 and 0.900 in units of the standardized similarity outcome. Revenue similarity is somewhat more predictable, with mean out-of-sample R^2 equal to 0.253, while expenditure similarity has mean out-of-sample R^2 equal to 0.160. These values should not be read as evidence that municipal budgets are mechanically determined by observables. Rather, they show that observable municipal characteristics contain meaningful information about expected fiscal similarity, while leaving substantial residual variation for the divergence measure.[11]

Because the divergence measure relies on the model as a benchmark, we also assess whether predicted similarity ranks municipality pairs in a meaningful way. Figure 4 reports a binned actual-versus-predicted plot, grouping observations by predicted budget similarity and plotting mean realized similarity against mean predicted similarity within each bin. The relationship is monotonic and close to the 45-degree line, indicating that the model preserves the ordering of municipality pairs well, although the fitted values are slightly compressed at the tails of the distribution. The figure therefore complements the out-of-sample R^2: R^2, MAE, and RMSE evaluate

[11] The visible step in average R^2 between 2007 and 2008 in Figure 3 coincides with the major change in the structure of Italian municipal budgets, which substantially increased the number of expenditure components included in our analysis (see Table 1). The richer post-2008 budget representation captures a larger share of systematic and predictable variation in fiscal policy, mechanically raising the model's overall explanatory power. Reassuringly, the model continues to perform well on ranking municipalities by predicted similarity, as shown in the binned actual-versus-predicted plot in Figure 4.

predictive accuracy, while the binned actual-versus-predicted plot evaluates whether predicted levels and rankings correspond to realized similarity.[12]

Having established that the benchmark has meaningful predictive content, we define fiscal divergence from the model's prediction errors. For each municipality pair in each region-year cell, we compute the absolute difference between observed and predicted similarity, $|r_{ij}| = |y_{ij} - \hat{y}_{ij}|$. This prediction error measures the gap between the pair's realized budget similarity and the level expected on the basis of observed characteristics. We then move from the pair-year to the municipality-year level by assigning each pairwise error to both municipalities in the pair and averaging the resulting absolute errors by municipality and year.

One concern is that prediction errors may reflect unobserved demographic change rather than unexpected policy behavior. This does not appear to drive the results: changes in census characteristics explain little of the average prediction error ($R^2 = 0.06$).[13] This supports our interpretation of the policy divergence index as capturing unexpected policy behavior rather than unobserved socio-demographic dynamics.

Appendix Figures 18 and 19 present the geographical distribution of policy divergence across Italian regions and provinces, with darker shades identifying the areas with higher values of aggregated fiscal divergence. We can see that there is variation across region and time, with the southern regions, along with Sardegna and Sicilia, experiencing rising divergence between 2005 and 2015. In Northern and Central Italy, the trend is more heterogeneous, without a clear direction. The province-level map further shows substantial within-region heterogeneity, particularly in the South and on the islands, suggesting that the regional aggregation in the headline panel masks meaningful within-region variation.

5 Fiscal divergence and electoral incentives

We then examine whether our divergence measure, which captures unexpected fiscal choices relative to peers, is affected by electoral incentives by testing for the presence of political budget cycles (PBCs).

PBCs arise from politicians' tendency to manipulate fiscal policies to gain electoral advantage. For instance, Nordhaus (1975) proposes that governments exploit a Phillips curve trade-off, implementing stimulatory policies before elections to attract

[12] For numerical stability, the XGBoost model is fitted on a standardized version of the cosine similarity, defined as the raw cosine divided by its within-region-year standard deviation. The values shown on the axes of Figure 4 and the prediction-error metrics in Table 3 are therefore expressed in this standardized scale rather than on the original [0, 1] cosine scale; the standardization is monotonic, so it does not affect the ordering of municipality pairs nor the interpretation of the actual-versus-predicted plot. In raw cosine units, the average within-region-year similarity is around 0.55-−0.61 and the corresponding standard deviation is roughly 0.13.

[13] Socio-demographic features come from the 2011 census and are held fixed across budget years, while budget data vary over 2000–2015. To assess whether this induces a mechanical demographic component in the prediction errors, we use the 2001 census wave and regress each municipality's average prediction error on standardized 2001–2011 changes in all socio-demographic variables. Figure 17 reports coefficients and confidence intervals for an illustrative subset. Most estimates are small and unsystematic.

votes from myopic voters. Similarly, Rogoff and Sibert (1988) argue that politicians signal their competence by increasing spending before elections. Building on this, Shi and Svensson (2006) further argue that all politicians, regardless of competence, may boost public goods provision before elections to influence uninformed voters, with stronger cycles observed in developing countries. Offering a different perspective, Drazen and Eslava (2010) suggest that in established democracies with informed voters, PBCs arise not from perceived competence but from politicians' manipulation of spending preferences. In this view, politicians may adjust the composition of spending to align with voter priorities, thereby signaling alignment with public interests. Although citizens may not have detailed knowledge of the specific budget breakdowns, they can still sense the impact of fiscal decisions through their experience of public goods and taxation.

In this framework, our divergence index serves as an indicator of the gap between actual budget outcomes and the budget composition predicted by local needs and characteristics. The literature therefore suggests that electoral incentives may affect expenditure and revenue differently. Expenditure choices are often visible and discretionary, so pre-electoral adjustments may bring spending composition closer to predictable local needs. Revenue-side responses are more instrument-specific: incumbents may selectively adjust tax rates, fees, transfers, or collection effort, which can generate departures from the revenue profile predicted by municipal characteristics (Alesina & Paradisi, 2017; Foremny & Riedel, 2014; Rogoff & Sibert, 1988). The analysis uses the staggered election schedules of over 8000 Italian municipalities. The staggered timing results from historical factors and from instances in which government crises led to the premature end of electoral mandates before their scheduled deadlines.

As suggested by previous research exploiting these features of the Italian local political system, exogenous timing helps separate the impact of the electoral cycle from the effects of events occurring in a given year and other confounders (Repetto, 2018). Figure 5 shows the number of elections by year.

Formally, we estimate the following equation:

$$\text{Divergence Index}_{it} = \alpha + \sum_{k=1}^{4} \beta_k \text{Electoral Cycle}_{i,t-k} + \mu_i + \epsilon_{it} \tag{6}$$

where $\text{Divergence Index}_{it}$ is the dependent variable representing the fiscal divergence index for municipality i at time t. The summation $\sum_{k=1}^{4} \beta_k \text{Electoral Cycle}_{i,t-k}$ includes indicator variables for each year of the electoral cycle from $k = 1$ (the year immediately after the election) to $k = 4$ (the year before the next election). The election year ($k = 0$) is omitted and serves as the reference category, so that the estimated coefficients β_k capture the difference in fiscal divergence between year k and the election year. We include a constant α, municipality fixed effects μ_i, and finally, ϵ_{it} represents the error term, which is clustered at the municipality level.[14]

[14] We do not include year fixed effects or year × region fixed effects, as the outcome variable has been constructed to inherently account for variations within years and across year × region interactions.

Our main results are illustrated in Figure 6, which displays the standardized estimated coefficients for each year of the electoral cycle relative to the election year ($k = 0$, reported as a hollow marker on the zero line), reported jointly for the full-budget, expenditure, and revenue divergence indices. The full-budget and expenditure series tell essentially the same story and show the clearest political-budget-cycle pattern: divergence is lower than in election years at $t-4$, $t-3$, and $t-1$, with standardized coefficients between roughly three and four percent of an outcome standard deviation. The $t-2$ coefficient is also negative and significantly different from zero at the 5% level for the full budget (-0.019, $p \approx 0.03$), although it is smaller in magnitude than the others and not statistically distinguishable from zero for expenditure alone. The closeness between the full-budget and expenditure series is also partly mechanical: in our 2008 reference matrix, the fiscal vector contains 4119 expenditure components and only 294 revenue components, so the cosine similarity over the full vector is dominated by the expenditure block. We discuss the implication for interpretation later in this section. The revenue series is qualitatively different: revenue divergence is positive and statistically significant at $t-4$ and $t-1$, close to zero at $t-3$, and negative but statistically indistinguishable from zero at $t-2$. This opposite sign should be interpreted cautiously, but it is consistent with the idea that electoral incentives operate through different fiscal margins. Expenditure cycles may reflect adjustments in visible spending composition (Drazen & Eslava, 2010; Ferraresi, 2020; Repetto, 2018), while revenue cycles may arise from selective changes in tax rates or other revenue instruments rather than broad convergence toward the revenue profile predicted by local characteristics (Alesina & Paradisi, 2017; Foremny & Riedel, 2014).

As a robustness check, we also re-estimate Equation (6) interacting each cycle indicator with a post-2008 dummy. The interaction coefficients generally move against the baseline cycle effects for both expenditure and revenue, pointing to post-

(a) Revenue

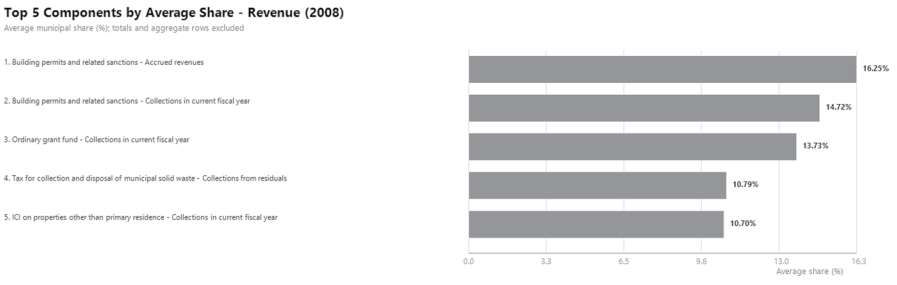

Fig. 7 Top 5 components by average share for revenue in 2008

(a) Commitments

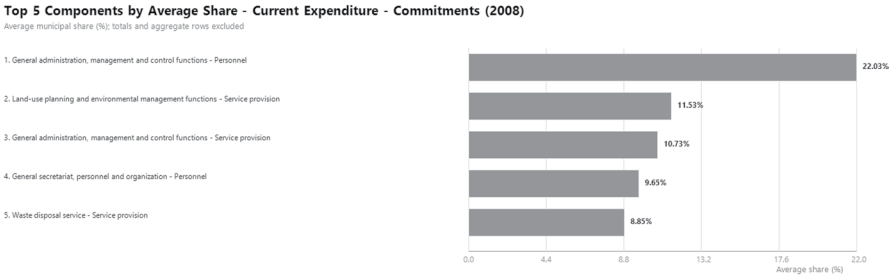

(b) Payments on current-year commitments

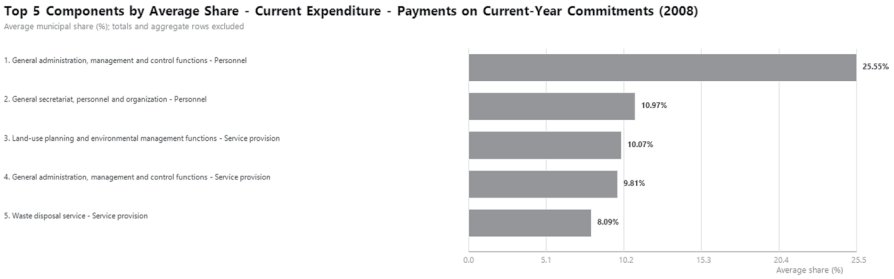

(c) Payments on prior-year residuals

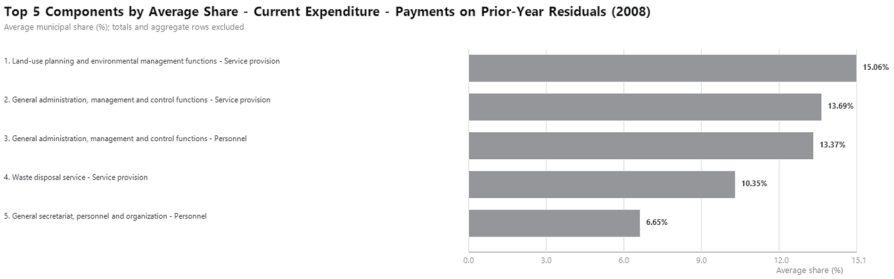

Fig. 8 Top 5 components by average share for current expenditure in 2008

2008 attenuation in line with Repetto (2018)'s evidence that the 2008 transparency reform reduced the scope for electoral-cycle behaviour at the local level.[15]

Our measure of fiscal divergence serves as a proxy for how closely budget composition adheres to the behavior predicted by local characteristics. The expenditure

[15] For the full-budget measure, the post-2008 interactions are positive across $t-4$ through $t-1$, offsetting the negative pre-electoral coefficients; the attenuation is strongest at $t-1$, where the implied post-2008 coefficient is close to zero. The expenditure interactions show the same offsetting pattern. On the revenue side, the significant interactions at $t-4$ and $t-2$ also move against the corresponding pre-2008 coefficients. We report all interaction coefficients in Appendix Figure 16. We also estimated versions of Equation (6) that weight municipality-year observations by the average out-of-sample R^2 of the corresponding region-year prediction model. The pattern remains very similar, indicating that the main result is not driven by low-quality prediction cells.

(a) Commitments

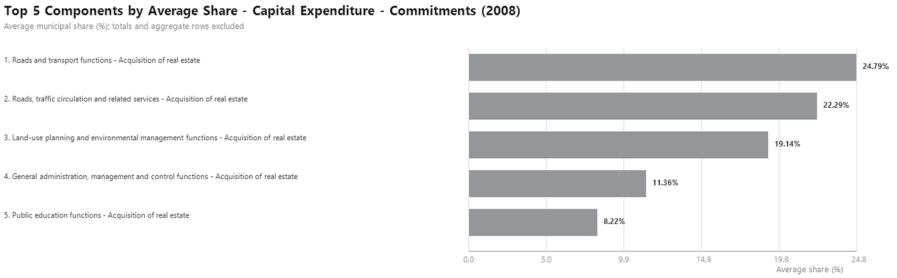

(b) Payments on current-year commitments

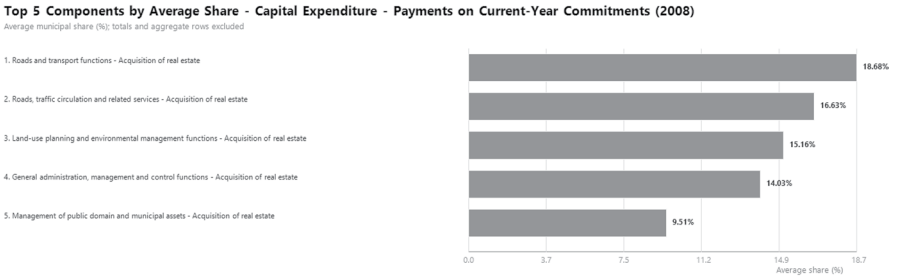

(c) Payments on prior-year residuals

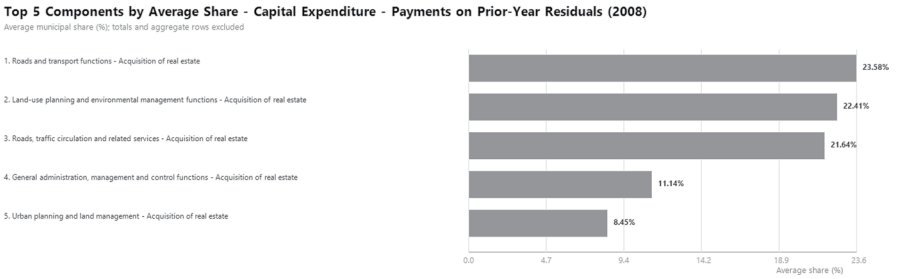

Fig. 9 Top 5 components by average share for capital expenditure in 2008

results are consistent with the view that, around elections, politicians avoid hard-to-explain departures from expected spending priorities and instead choose spending mixes that are more in line with predictable local needs. This interpretation fits the mechanism proposed by Drazen and Eslava (2010), where politicians adjust the composition of public spending to signal alignment with voters, and is also consistent with evidence that politicians are less likely to pursue risky policy experimentation close to elections (Bernecker et al., 2021). The revenue results point to a different margin: rather than converging toward the predicted revenue mix, municipalities appear to deviate from it in selected pre-electoral years. This is consistent with revenue-side political budget cycles being narrower and more instrument-specific than expenditure cycles.

Table 1 Budget Components used in Cosine Similarity Computation for the Years 2000–2015

Year	Number of Revenue Components	Number of Expense Components
2000	251	2331
2001	263	2331
2002	266	2331
2003	261	2331
2004	261	2331
2005	261	2331
2006	278	2331
2007	271	2394
2008	294	4119
2009	315	4155
2010	315	4185
2011	330	4185
2012	324	4185
2013	339	4245
2014	345	4205
2015	345	4205

Note: The table reports the number of budget variables used to calculate fiscal-similarity in policymaking across municipalities by year and type of component. "Expense components" aggregates the variables drawn from *quadro 4* (current expenditure) and *quadro 5* (capital expenditure) across the three accounting flows (impegni, pagamenti in conto competenza, pagamenti in conto residui). For 2008, these counts coincide with the breakdown discussed in the main text and in Appendix A (294 revenue, 1,977 current expenditure, and 2,142 capital expenditure components, for a total of 4,413)

6 Conclusion

In this paper, we developed a novel approach to measuring fiscal-policy divergence at the local level using comprehensive municipal budget data. By analyzing budget similarities across roughly 8000 Italian municipalities from 2000 to 2015, we introduced a measure of policy divergence as absolute out-of-sample prediction errors from flexible models of fiscal similarity among municipalities with similar socio-demographic profiles. Our method provides a clear and practical metric for capturing when local governments engage in unexpected policy actions. We validated our measure by demonstrating its sensitivity to political incentives in an empirical application to the political budget cycle. The evidence is strongest on the expenditure side: expenditure divergence is lower in several pre-electoral years than in election years, while revenue divergence follows a different and less stable pattern. The full-budget result, reported together with expenditure and revenue in the main analysis, closely follows the expenditure pattern.

In sum, our paper makes three key contributions. First, we develop a data-driven measure of fiscal policy divergence that captures the extent to which local governments' budget allocations deviate from expectations based on socio-demographic characteristics. This approach allows us to analyze the entire composition of fiscal policy, rather than focusing on individual budget categories in isolation. Second, we show that our method provides a flexible and consistent benchmark for cross-jurisdictional comparisons, allowing researchers to detect and quantify unexpected fiscal behaviors across a wide variety of contexts. Finally, through an application to Italian

Table 2 List of Socio-Demographic and Mayor Variables

Variable name	Variable name	Variable name
Population	Male Labor Participation	Density
Incidence of homeowners	Female Labor Participation	Population Under 6
Housing age	Labor Activity Rate	Population Over 75
Living area (sq m)	Municipal Unemployment	Old Age Dependency
Housing crowding	Functional Unemployment	Divorce Rate
Family size	Workforce Renewal Rate	Foreign Nationals
Elderly living alone	Agricultural Employment	Foreign Minors
Illiteracy rate	Manufacturing Employment	Foreign Employment
Education 15–19	Trade Employment	Social Vulnerability
Commuting for work	High-Skill Employment	Avg. Size Household
Commuting for study	Manual-Labor Employment	Internal Migration Net
Private mobility	Low-Skill Employment	Male International Migration
Public mobility	Job-Related Mobility	Female International Migration
No city center	Study-Related Mobility	Active/Inactive Youths Ratio
Urban density	Private Trans. Mobility	Employed Italians/Foreigners Ratio
Building density	Public Trans. Mobility	Housing Price Index
Young families	School Dropouts	Mayor Age
School dropout rate	Building Age	Mayor Gender
Migration population	Migration Segregation	Mayor Education
Employed immigrants	Income Below Poverty	Mayor Party Affiliation
Income €0–€10k	Income €10k–€15k	
Income €15k–€26k	Income €26k–€55k	
Income €55k–€75k	Income €75k–€120k	
Income above €120k		

Note: The table reports the municipal and mayor variables used as predictors of fiscal-policy similarity. The first sixty entries are the socio-demographic and economic variables drawn from the 2011 Italian decennial census and treated as time-invariant. The four entries at the end of the right-hand column (Mayor Age, Mayor Gender, Mayor Education, Mayor Party Affiliation) are extracted from the *Anagrafe degli Amministratori Locali e Regionali* of the Ministry of Interior and vary across mayoral terms. The same variable names are used in the validation regressions of Equation (4) (Figures 2, 10, 11, 12, 13) and in the XGBoost feature-importance tables (Tables 4, 5, 6). Income brackets are reported in euros

Table 3 Out-of-sample prediction performance by budget measure

Outcome	Mean R^2	Median R^2	Mean MAE	Mean RMSE	Mean MSE	Cells
Full budget	0.189	0.188	0.718	0.900	0.811	1,520
Expenditure	0.160	0.154	0.732	0.916	0.840	1,520
Revenue	0.253	0.247	0.694	0.863	0.747	1,520

Notes: The table reports out-of-sample prediction performance for the region-year XGBoost models. R^2, mean absolute error (MAE), root mean squared error (RMSE), and mean squared error (MSE) are computed over region-year-fold cells from the 5-fold cross-validation procedure. The prediction target is the standardized cosine similarity (raw cosine divided by its within-region-year standard deviation), so MAE, RMSE, and MSE are in units of this standardized outcome rather than in raw cosine points. "Cells" reports the number of fold-level estimates (16 years × 19 regions × 5 folds = 1,520)

municipalities, we demonstrate the empirical relevance of our measure by showing that electoral incentives are reflected mainly in expenditure divergence, shed-

Table 4 Top 10 XGBoost predictors: full budget similarity

Rank	Feature	Share of total gain
1	Difference in Commuting for Study	6.89%
2	Difference in Agricultural Employment	6.75%
3	Difference in Foreign Employment	6.25%
4	Difference in Income Above €120k	5.38%
5	Difference in Income Below Poverty	4.67%
6	Difference in Elderly Living Alone	3.75%
7	Difference in Education 15–19	3.52%
8	Difference in Foreign Minors	2.97%
9	Difference in Active/Inactive Youths Ratio	2.71%
10	Difference in Male International Migration	2.55%

Notes: The table reports the ten predictors with the largest total gain in the full-budget XGBoost prediction model. "Difference in *X*" refers to the standardized absolute pairwise difference in variable *X* between the two municipalities in a pair. Shares are computed as each feature's total gain divided by total gain across all predictors

Table 5 Top 10 XGBoost predictors: expenditure similarity

Rank	Feature	Share of total gain
1	Difference in Agricultural Employment	7.61%
2	Difference in Commuting for Study	7.04%
3	Difference in Income Above €120k	6.51%
4	Difference in Income Below Poverty	5.34%
5	Difference in Education 15–19	4.03%
6	Difference in Workforce Renewal Rate	3.28%
7	Difference in Foreign Employment	3.14%
8	Difference in Male International Migration	2.72%
9	Difference in Illiteracy Rate	2.49%
10	Difference in Income €75k–€120k	2.39%

Notes: The table reports the ten predictors with the largest total gain in the expenditure-side XGBoost prediction model. Shares are computed as each feature's total gain divided by total gain across all predictors

ding light on how politicians adjust the composition of local policymaking around elections.

Importantly, our measure is descriptive and does not impose a normative or welfare interpretation on fiscal divergence. Deviations from predicted policies may reflect innovative policymaking, but they may also capture inefficiencies or other forms of misallocation. The interpretation of divergence, therefore, depends on the specific institutional context and research question at hand.

The core contribution of our study is methodological rather than country-specific. We propose a general framework for measuring policy divergence that can be implemented in any institutional context where detailed budget and socio-demographic

Table 6 Top 10 XGBoost predictors: revenue similarity

Rank	Feature	Share of total gain
1	Difference in Foreign Employment	13.31%
2	Difference in Elderly Living Alone	6.49%
3	Difference in Family Size	4.53%
4	Difference in Employed Italians/Foreigners Ratio	4.32%
5	Difference in Foreign Minors	4.21%
6	Difference in Active/Inactive Youths Ratio	4.10%
7	Difference in Population Over 75	4.06%
8	Difference in Housing Price Index	3.28%
9	Difference in Income Above €120k	2.89%
10	Difference in Trade Employment	2.59%

Notes: The table reports the ten predictors with the largest total gain in the revenue-side XGBoost prediction model. Shares are computed as each feature's total gain divided by total gain across all predictors

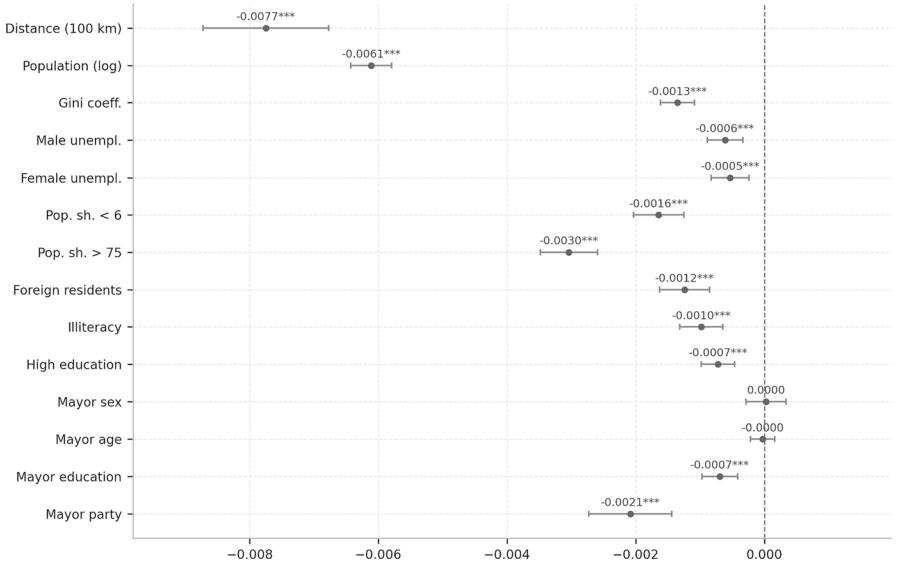

Fig. 10 Validation Policy Similarity - Excluding less discretionary budget items. Note: This graph shows the coefficients and confidence intervals of various municipal characteristics differences between a pair of municipalities (standardized) on their fiscal policy similarity. The cosine similarity is recomputed after excluding budget components that are typically less discretionary at the municipal level. The excluded categories are listed in Appendix A.2: on the revenue side, intergovernmental transfers from the State, regions, other public bodies, and the European Union, as well as transfers earmarked for personnel costs and for capital investment; on the expenditure side, social, education, and welfare items (e.g., *servizi per l'infanzia e i minori*, *strutture residenziali e di ricovero per anziani*, *assistenza scolastica*, *trasferimenti*) in both quadro 4 and quadro 5. The stars denote significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The x-axis shows the magnitude of the coefficients, while the y-axis lists the variables. The error bars represent the 95% confidence intervals for each coefficient

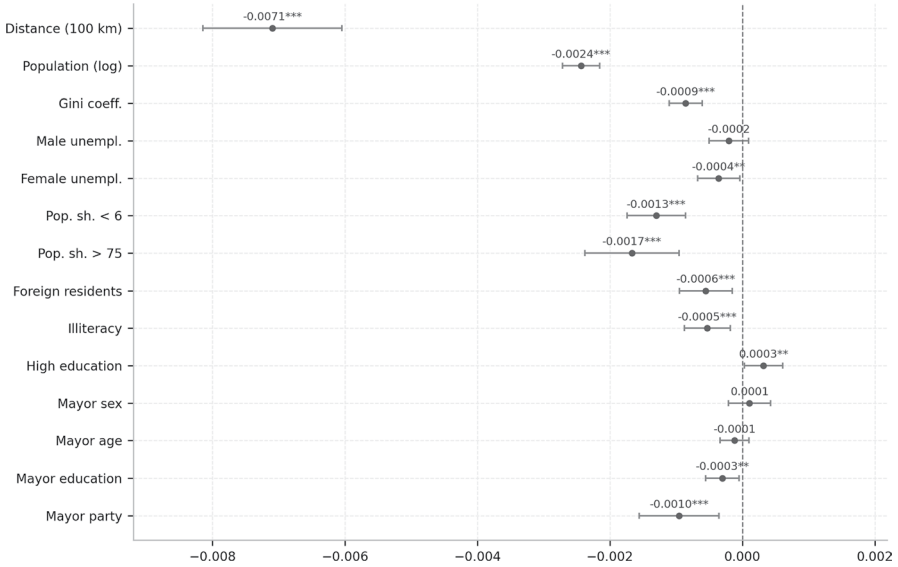

Fig. 11 Validation Policy Similarity - Full Set of Regressors. Note: This graph shows the coefficients and confidence intervals of selected regressors from a regression of Equation (4) that includes the complete set of differences in municipal and mayor characteristics listed in Table 2. The stars denote significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The x-axis shows the magnitude of the coefficients, while the y-axis lists the variables. The error bars represent the 95% confidence intervals for each coefficient

data are available. Its external validity therefore hinges primarily on data availability and on the degree of fiscal discretion enjoyed by subnational governments; in settings where local authorities retain meaningful autonomy over expenditure composition, we would expect qualitatively similar patterns to emerge, even if the magnitude of the effects varies across institutional environments.

The proposed framework lends itself to a wide range of applications beyond the electoral cycle. Future research could use our measure to study policy innovation and diffusion across jurisdictions, the effects of political alignment or decentralization reforms, the relationship between unexpected fiscal choices and outcomes such as corruption or citizen satisfaction, or the fiscal consequences of the rise of populist parties at the local level.

Budget structure

This appendix describes, in plain terms, how municipal budget records are converted into the fiscal vector used for similarity analysis.

The raw accounting records are organized along four dimensions: *quadro* (major budget block), *subquadro* (accounting flow block), *voce* (row item), and *colonna* (column item). In the working dataset, each component is stored as a variable named `vQQ_VVV_C`, where `QQ` is the quadro code, `VVV` is the voce code, and `C` is the col-

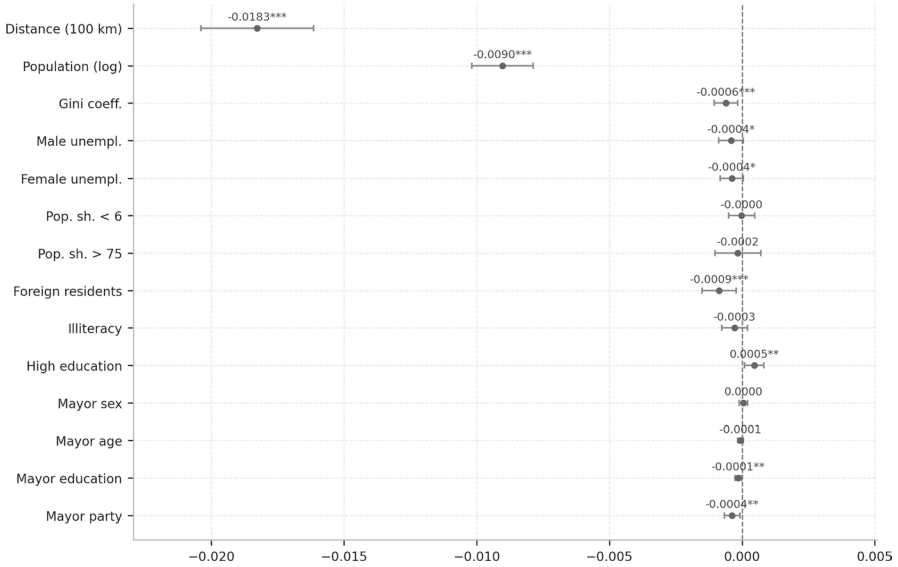

Fig. 12 Validation Policy Similarity - Quadratic Specification. Note: This figure replicates the main validation exercise (Equation 4) augmented with quadratic terms in all socio-demographic and mayor covariates listed in Table 2. Only the linear-term coefficients for the same set of regressors shown in Figure 2 are displayed. The stars denote significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The 95% confidence intervals are based on multi-way clustered standard errors

umn code. A single variable therefore corresponds to one specific budget cell, namely one row-column intersection inside a given block.

For the fiscal similarity index, the vector is built from three budget blocks: Quadro 02 (revenues), Quadro 04 (current expenditure), and Quadro 05 (capital expenditure). In the 2008 restricted matrix used here as reference, this yields 4,413 components in total. Of these, 294 come from revenues, 1,977 from current expenditure, and 2,142 from capital expenditure.

A key institutional detail concerns subquadri. Quadro 02 has no A/B/C split, while Quadri 04 and 05 are each partitioned into three subquadri: A (*impegni*), B (*pagamenti in conto competenza*), and C (*pagamenti in conto residui*). In 2008, this split is perfectly balanced in dimensional terms: Quadro 04 contributes 659 components in each of A, B, and C, and Quadro 05 contributes 714 components in each of A, B, and C. In practice, subquadri are encoded by distinct voce code ranges (A: 005–188, B: 190–373, C: 375–560), which is why similarly named rows can appear at different voce codes.

The support of the data is almost, but not fully, Cartesian in rows and columns. Revenues use 3 columns and contain 98 observed components per column (for 294 in total). In Quadro 04, each subquadro has 66 components in columns 1–9 and 65 in column 10. In Quadro 05, each subquadro has 65 components in columns 1–10 and 64 in column 11. Hence, each element of the fiscal vector is still a well-defined budget cell, but the final feature space reflects the empirical reporting support rather than an imposed complete grid.

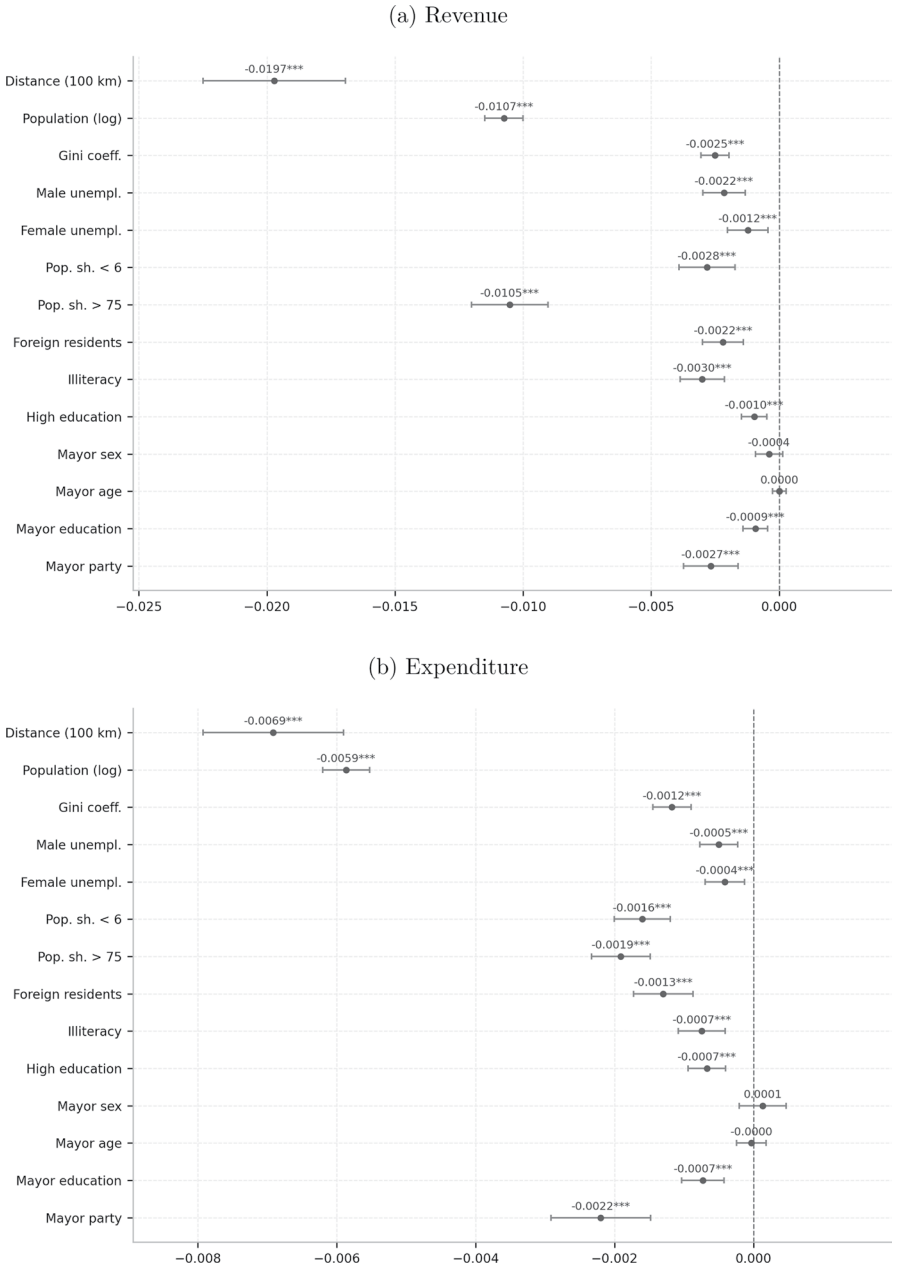

Fig. 13 Validation Policy Similarity by Budget Block. Note: This figure replicates the validation exercise of Equation (4) separately for revenue similarity (panel a) and expenditure similarity (panel b). The set of regressors and the conventions for stars and confidence intervals are the same as in Figure 2.

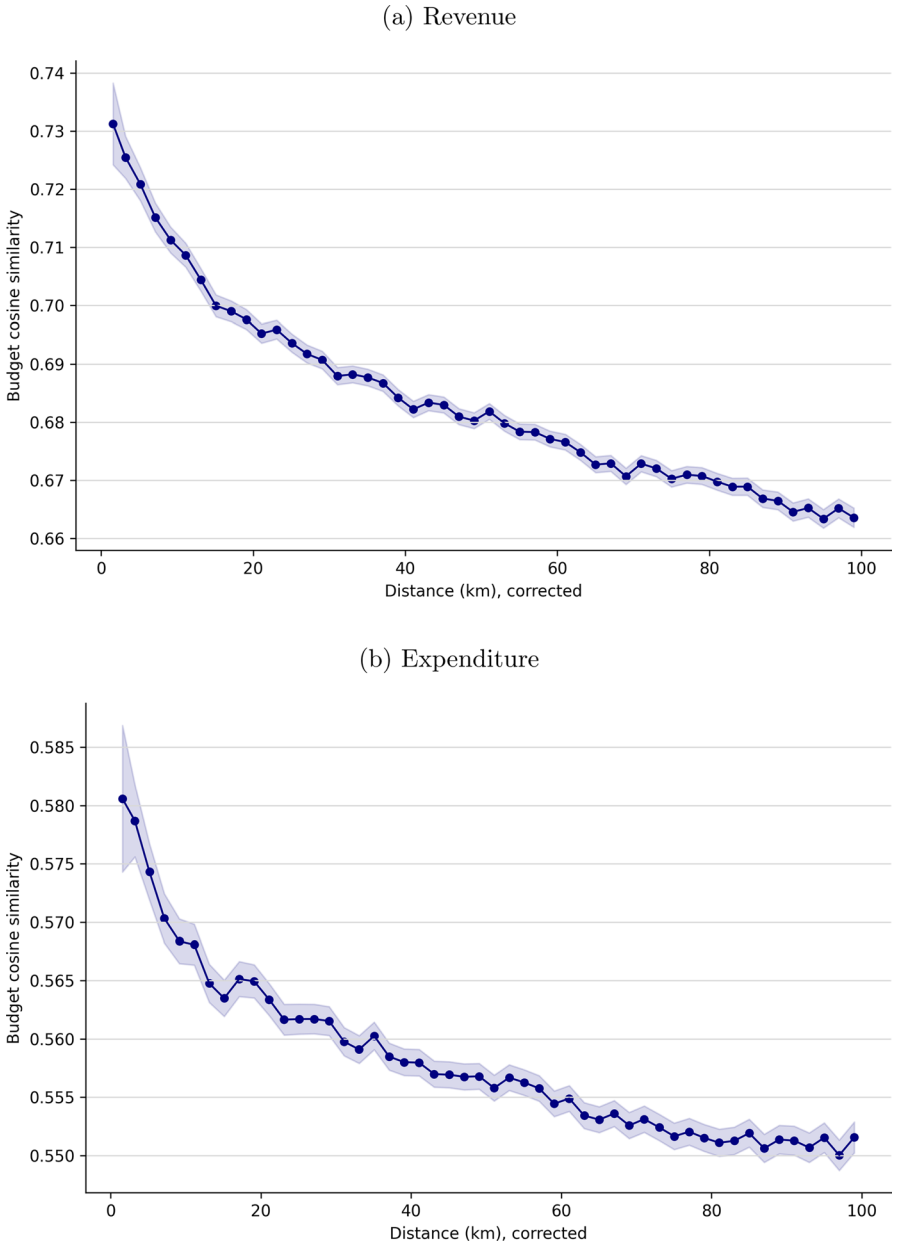

Fig. 14 Budget Similarity and Geographical Distance by Budget Block. Note: The binscatter plots illustrate the relationship between budget similarity and geographical distance separately for revenue and expenditure.

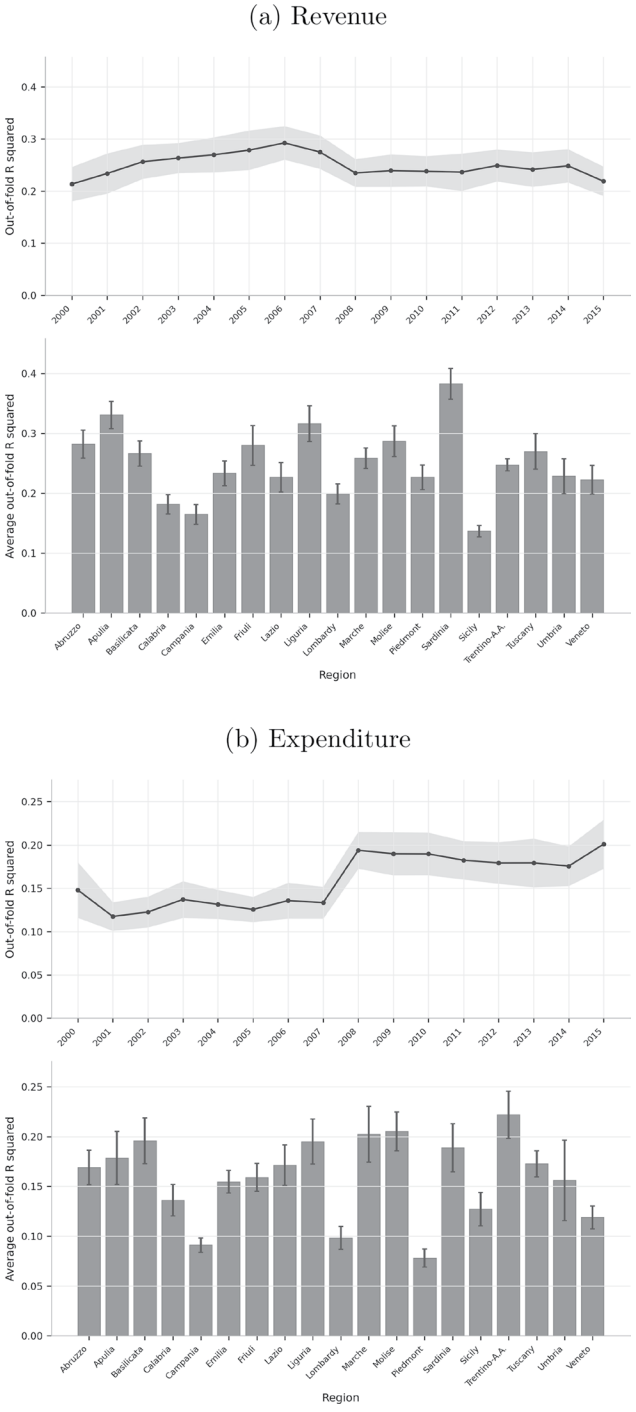

Fig. 15 Out-of-Sample Prediction Performance by Budget Block. Note: These figures report the temporal and regional variation in out-of-sample R^2 for the revenue (panel a) and expenditure (panel b) XGBoost models. Panel structure and shaded confidence intervals follow Figure 3.

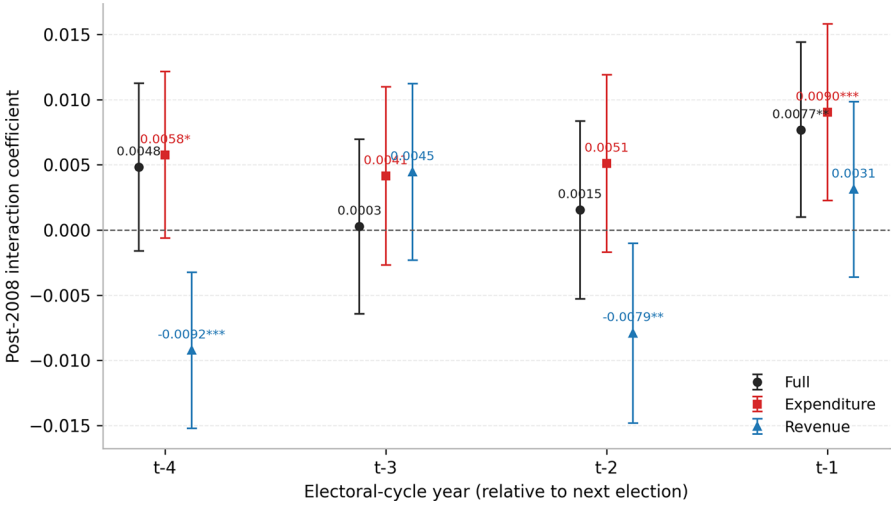

Fig. 16 Political Incentives and Fiscal Divergence: Post-2008 Interaction. Note: The figure reports the post-2008 × electoral-cycle interaction coefficients from a re-estimation of Equation (6) that adds an indicator for years $\geq$ 2008 and its interaction with the four electoral-cycle dummies. Coefficients are shown for the full-budget, expenditure, and revenue divergence indices. The omitted category is the election year. Standard errors are clustered at the municipality level. Numbers next to each marker report the point estimate followed by significance stars (*** p<0.01, ** p<0.05, * p<0.10)

For each municipality-year, each included cell is transformed into a share by dividing the raw amount by the relevant total for its block. The fiscal profile is then the stacked vector of these normalized shares. Fiscal similarity between municipalities is computed as cosine similarity between these vectors, so dimensional comparability is determined by the shared component structure described above.

Discussion of the main components

Figure 7 shows that the revenue structure in 2008 is mainly driven by building-related revenues, intergovernmental transfers, and local taxes. The largest component is *building permits and related sanctions – accrued revenues* (16.25%), followed by *building permits and related sanctions – collections in current fiscal year* (14.72%). This suggests that construction activity and the related administrative charges are an important source of revenue. Another major item is the *ordinary grant fund – collections in current fiscal year* (13.73%), which highlights the role of transfers from higher levels of government. Revenue is also supported by local taxation, especially the *tax for collection and disposal of municipal solid waste* (10.79%) and the *ICI on properties other than primary residence* (10.70%).

Current expenditure is dominated by personnel and service provision activities. In Figure 8(a), the largest component is *general administration, management and control functions – personnel* (22.03%), indicating that administrative staffing absorbs a substantial share of current spending. Other relevant items are *land-use planning and environmental management functions – service provision* (11.53%), *general administration, management and control functions – service provision* (10.73%), and *waste*

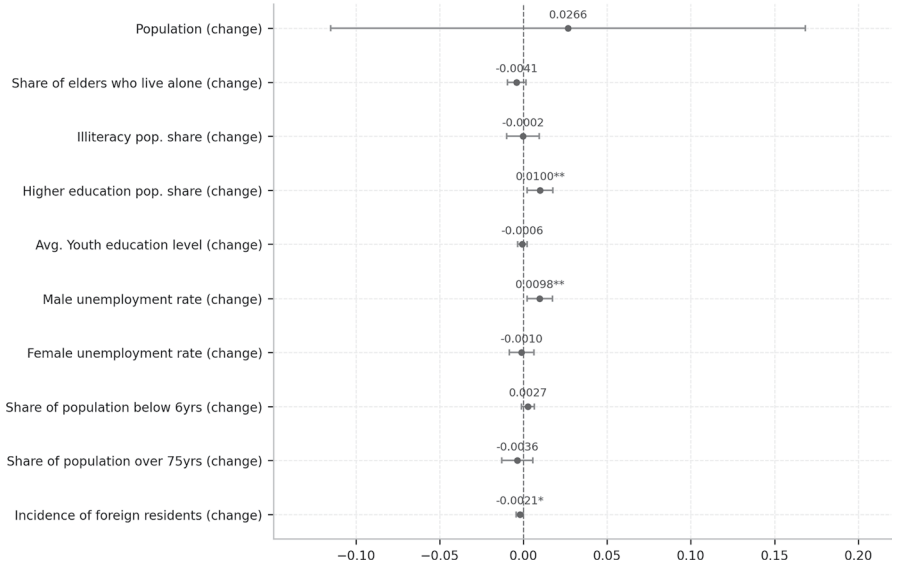

Fig. 17 Correlation between Residuals and Change in Socio-demographic Characteristics. Note: This graph shows the coefficients and confidence intervals of selected regressors from a regression of each municipality's average absolute prediction error (from Equation (5)) on the standardized 2001–2011 change in census socio-demographic characteristics. Each regressor is the absolute change between the two census waves rescaled by the cross-sectional standard deviation of the 2001 vintage. The stars denote significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The x-axis shows the magnitude of the coefficients, while the y-axis lists the variables. The error bars represent the 95% confidence intervals for each coefficient

disposal service – service provision (8.85%). The same pattern appears in Figure 8(b), where *general administration, management and control functions – personnel* rises to 25.55%. This confirms that wages and routine administrative functions are central in current municipal expenditure. In Figure 8(c), the most important items are more service-oriented, especially *land-use planning and environmental management functions – service provision* (15.06%) and *general administration, management and control functions – service provision* (13.69%), suggesting that outstanding payments from previous years are often linked to service contracts and operating obligations.

Capital expenditure is highly concentrated in infrastructure and land-related investment. In Figure 9(a), the largest components are *roads and transport functions – acquisition of real estate* (24.79%), *roads, traffic circulation and related services – acquisition of real estate* (22.29%), and *land-use planning and environmental management functions – acquisition of real estate* (19.14%). This indicates that capital spending is primarily directed toward transport infrastructure and land management. A similar composition appears in Figure 9(b) and Figure 9(c), where roads, traffic, and land-use planning remain the dominant items. In particular, prior-year residual payments are strongly concentrated in *roads and transport functions – acquisition of real estate* (23.58%), *land-use planning and environmental management functions – acquisition of real estate* (22.41%), and *roads, traffic circulation and related services – acquisition of real estate* (21.64%). Overall, the evidence suggests that capital

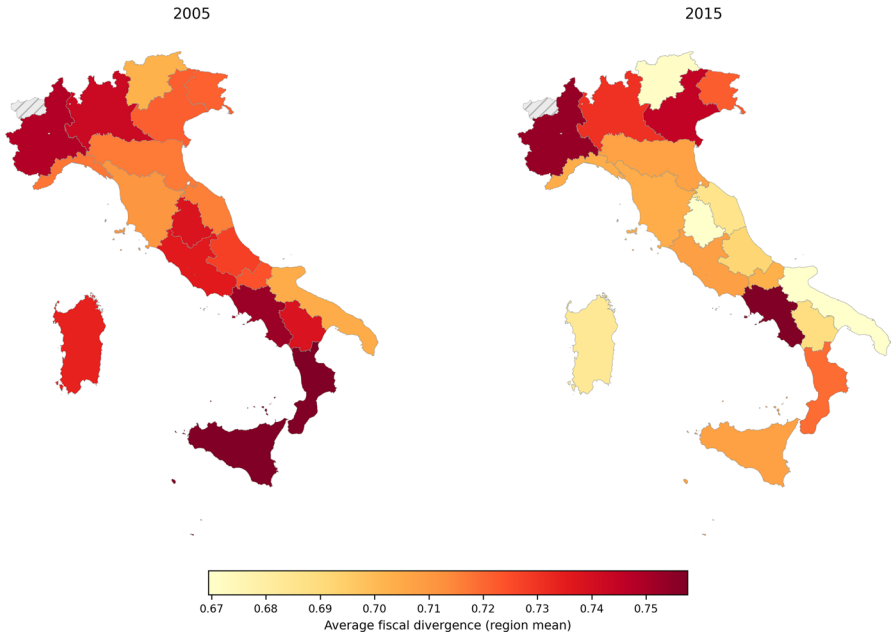

Fig. 18 Geographical Distribution of Fiscal Policy Divergence – Region Level. *Notes:* The figure reports the geographical distribution of fiscal-policy divergence in 2005 and 2015. For each region, divergence is computed as the average of the municipality-level mean absolute prediction error from the XGBoost model in Equation (5). Darker shades indicate higher divergence; lighter shades indicate lower divergence. The two panels share the same colour scale, so values are directly comparable across years. Valle d'Aosta is left blank because the within-region pair structure cannot be built (only one municipality enters our estimation sample)

expenditure follows a multi-year implementation profile, with infrastructure-related commitments generating payments over more than one fiscal year.

Less discretionary budget items excluded in the robustness analysis

In the robustness exercise reported in Figure 10 we re-compute the cosine similarity after dropping budget components that, in the Italian institutional setting, are largely formula-based or otherwise outside the direct control of the municipality. The selection is conservative: we remove items whose magnitude is mechanically tied to higher-level rules, transfers, or earmarked funding rather than to discretionary local decisions.

On the revenue side (*quadro 2*), we exclude the following *voci*:

- Ordinary, equalization, and consolidated grant funds from the central government;

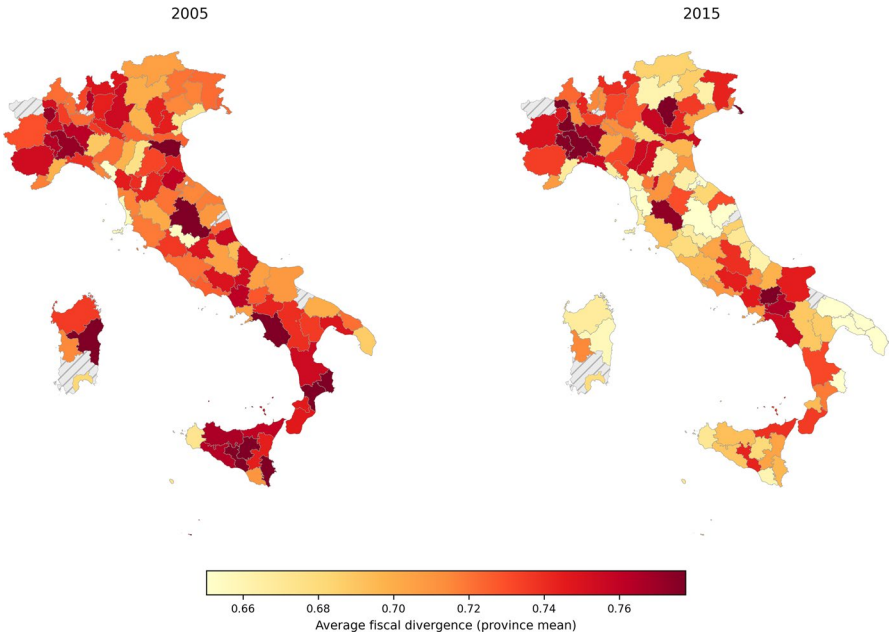

Fig. 19 Geographical Distribution of Fiscal Policy Divergence – Province Level. *Notes:* The figure reports the same evidence as Figure 18 aggregated at the province level. For each province, divergence is the average of the municipality-level mean absolute prediction error from the XGBoost model in Equation (5). Darker shades indicate higher divergence; lighter shades indicate lower divergence. The two panels share the same colour scale, so values are directly comparable across years. Provinces fully covered by Valle d'Aosta or with no municipalities in the estimation sample are left blank

- Transfers earmarked for personnel and contractual increases;
- Transfers from the State for specific functions and from regions for delegated functions;
- Transfers from the European Union and from other public bodies;
- Transfers from the State and the regions earmarked for capital investment.On the expenditure side (*quadri 4 and 5*), we exclude the same set of *voci* across the three accounting flows (impegni, pagamenti in conto competenza, pagamenti in conto residui):

- Childcare and services for children and minors (*servizi per l'infanzia e i minori*);
- Residential and care facilities for the elderly (*strutture residenziali e di ricovero per anziani*);
- School assistance, transport, canteens, and related school services (*assistenza scolastica*, *trasporti*, *refezione*, *altri servizi*);
- Transfers (*trasferimenti*) used to channel earmarked higher-level funds.These items account for a non-trivial share of the budget but, in our setting, are typically allocated through national or regional formulas linked to demographic indicators

(population, number of children, number of elderly, share of unemployed), and are therefore precisely the components for which a mechanical link with local characteristics is most plausible.

Author contributions All authors contributed equally to each part of the research.

Funding Open access funding provided by Università degli Studi di Roma La Sapienza within the CRUI-CARE Agreement.

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

References

Agrawal, D. R., Breuillé, M.-L., & Le Gallo, J. (2025). Tax competition with intermunicipal cooperation. *National Tax Journal, 78*(1), 5–43.

Agrawal, D. R., Hoyt, W. H., & Wilson, J. D. (2022). Local policy choice: Theory and empirics. *Journal of Economic Literature, 60*(4), 1378–1455.

Akhmedov, A., & Zhuravskaya, E. (2004). Opportunistic political cycles: Test in a young democracy setting. *The Quarterly Journal of Economics, 119*(4), 1301–1338.

Alesina, A., & Paradisi, M. (2017). Political budget cycles: Evidence from Italian cities. *Economics & Politics, 29*(2), 157–177.

Alesina, A., Cassidy, T., & Troiano, U. (2019). Old and young politicians. *Economica, 86*(344), 689–727.

Allers, M., & Paul Elhorst, J. (2005). Tax mimicking and yardstick competition among local governments in the Netherlands. *International Tax and Public Finance, 12*(4), 493–513.

Anderson, J. E., & Van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American economic review, 93*(1), 170–192.

Ash, E., & Galletta, S. (2023). How cable news reshaped local government. *American Economic Journal: Applied Economics, 15*(4), 292–320.

Ash, E., & Galletta, S. (2024). The Partisan Divergence in U.S. Local Policy, 1972–2017. *SSRN*. ETH Zürich.

Ash, E., Galletta, S., & Giommoni, T. (2025). A machine learning approach to analyze and support anticorruption policy. *American Economic Journal: Economic Policy, 17*(2), 162–193.

Bandiera, O., Prat, A., & Valletti, T. (2009). Active and passive waste in government spending: Evidence from a policy experiment. *American Economic Review, 99*(4), 1278–1308.

Baskaran, T., & Hessami, Z. (2025). Women in political bodies as policymakers. *Review of Economics and Statistics, 107*(6), 1501–1517.

Baskaran, T., Hessami, Z., & Schirner, S. (2024). Young versus old politicians and public spending priorities. *Journal of Economic Behavior & Organization, 225*, 88–106.
Bernecker, A., Boyer, P. C., & Gathmann, C. (2021). The role of electoral incentives for policy innovation: Evidence from the us welfare reform. *American Economic Journal: Economic Policy, 13*(2), 26–57.
Besley, T., & Case, A. (1995). Incumbent behavior: Vote seeking, tax setting and yardstick competition. *American Economic Review, 85*(1), 25–45.
Bochenkova, A., Buonanno, P., & Galletta, S. (2023). Fighting violence against women: The role of female political representation. *Journal of Development Economics, 164*, Article 103140.
Bordignon, M., Cerniglia, F., & Revelli, F. (2003). In search of yardstick competition: A spatial analysis of italian municipality property tax setting. *Journal of Urban Economics, 54*(2), 199–217.
Bosch, N., & Solé-Ollé, A. (2007). Yardstick competition and the political costs of raising taxes: An empirical analysis of spanish municipalities. *International Tax and Public Finance, 14*(1), 71–92.
Brollo, F., & Troiano, U. (2016). What happens when a woman wins an election? Evidence from close races in Brazil. *Journal of Development Economics, 122*, 28–45.
Brueckner, J. K. (1998). Testing for strategic interaction among local governments: The case of growth controls. *Journal of Urban Economics, 44*(3), 438–467.
Buettner, T. (2001). Local business taxation and competition for capital: The choice of the tax rate. *Regional Science and Urban Economics, 31*(2–3), 215–245.
Chen, T., & Guestrin, C. (2016). *XGBoost: A scalable tree boosting system* (pp. 785–794)
Daniele, G., & Giommoni, T. (2025). Fiscal rules, corruption, and electoral accountability. *The Journal of Politics, 87*(3), 857–871.
Drazen, A., & Eslava, M. (2010). Electoral manipulation via voter-friendly spending: Theory and evidence. *Journal of Development Economics, 92*(1), 39–52.
Feler, L., & Senses, M. Z. (2017). Trade shocks and the provision of local public goods. *American Economic Journal: Economic Policy, 9*(4), 101–143.
Ferraresi, M. (2020). Political cycles, spatial interactions and yardstick competition: Evidence from Italian cities. *Journal of Economic Geography, 20*(4), 1093–1115.
Ferraz, C., & Finan, F. (2009). Motivating politicians: The impacts of monetary incentives on quality and performance. National Bureau of Economic Research.
Ferreira, F., & Gyourko, J. (2014). Does gender matter for political leadership? The case of US mayors. *Journal of Public Economics, 112*, 24–39.
Foremny, D., & Riedel, N. (2014). Business taxes and the electoral cycle. *Journal of Public Economics, 115*, 48–61.
Gagliarducci, S., & Nannicini, T. (2013). Do better paid politicians perform better? Disentangling incentives from selection. *Journal of the European Economic Association, 11*(2), 369–398.
Galletta, S. (2020). Direct democracy, partial decentralization and voter information: Evidence from Swiss municipalities: S. Galletta. *International Tax and Public Finance, 27*(5), 1174–1197.
Grembi, V., Nannicini, T., & Troiano, U. (2016). *Do fiscal rules matter?, American Economic Journal: Applied Economics* pp. 1–30
Helm, I., & Stuhler, J. (2024). The dynamic response of municipal budgets to revenue shocks. *American Economic Journal: Applied Economics, 16*(4), 484–527.
Janeba, E., & Osterloh, S. (2013). Tax and the City–A Theory of Local Tax Competition. *Journal of Public Economics, 106*, 89–100.
Janeba, E., Köhler, F., Rostam-Afschar, D., & Steger, P. (2025). The Fiscal and Economic Effects of the BioNTech Shock Conference Paper.
Martínez, L. R. (2025). Natural Resource Rents, Local Taxes, and Government Performance: Evidence from Colombia. Review of Economics and Statistics.
Nordhaus, W. D. (1975). The political business cycle. *The Review of Economic Studies, 42*(2), 169–190.
Oates, W. E. (1972). Fiscal federalism. Books.
Repetto, L. (2018). Political budget cycles with informed voters: Evidence from Italy. *The Economic Journal, 128*(616), 3320–3353.
Revelli, F. (2001). Spatial patterns in local taxation: Tax mimicking or error mimicking? *Applied Economics, 33*(9), 1101–1107.
Revelli, F. (2007). Local government efficiency patterns in norway. *Journal of Urban Economics, 62*(1), 123–145.

Rogoff, K., & Sibert, A. (1988). Elections and macroeconomic policy cycles. *The Review of Economic Studies, 55*(1), 1–16.

Shi, M., & Svensson, J. (2006). Political budget cycles: Do they differ across countries and why? *Journal of Public Economics, 90*(8–9), 1367–1389.

Solé-Ollé, A. (2006). Expenditure spillovers and fiscal interactions: Empirical evidence from local governments in spain. *Journal of Urban Economics, 59*(1), 32–53.