

Efficient estimation of large spatially-clustered spatial autoregressive models for the analysis of spillovers in the voting concentration of Italian municipalities[☆]

Roy Cerqueti ^{a,b}, Valerio Ficcadenti ^c, Paolo Maranzano ^{d,e}, Raffaele Mattera ^f*

^a Department of Social and Economic Sciences, Sapienza University of Rome, Rome, Italy

^b GRANEM, University of Angers, SFR CONFLUENCES, F-49000, Angers, France

^c School of Business, London South Bank University, London, United Kingdom

^d Department of Economics, Management and Statistics, University of Milano-Bicocca, Piazza dell'Ateneo Nuovo, n.1, 20126, Milan, Italy

^e Fondazione Eni Enrico Mattei, Corso Magenta, n.63, 20123, Milan, Italy

^f Department of Mathematics and Physics, University of Campania "Luigi Vanvitelli", Caserta, Italy

ARTICLE INFO

Keywords:

Clusterwise regression
Spatial clustering
Spatial regimes
Gini and Shannon indexes
Italian Municipal data
Spatial statistics for electoral studies

ABSTRACT

This paper examines spatial spillovers in voting concentration across Italian municipalities, a key aspect for understanding electoral competition and political representation. To estimate heterogeneous spatial dependencies at the cluster level, while accounting for a broad set of socio-economic and institutional covariates, we adopt the recently introduced spatially-clustered spatial autoregressive (SCSAR) model. The standard SCSAR requires the evaluation of log-determinants of matrices whose dimension grows with the number of spatial units. Due to the large number of Italian municipalities, we address the computational burden of maximizing the log-likelihood function in SCSAR by means of matrix factorization methods. A computational benchmarking experiment shows that sparse factorization substantially improves scalability relative to the eigenvalue-based implementation, making SCSAR feasible for large municipal systems. This approach therefore enables a fast identification of clusters of Italian municipalities characterized by similar structural relationships. We find that global SAR models provide limited explanatory power of voting concentration compared to clustered specifications. Moreover, spatial spillovers are positive in the main specifications but vary substantially across clusters and electoral chambers, confirming that a single global spatial autoregressive coefficient masks important territorial heterogeneity. Economic and demographic variables are the most robust predictors of voting concentration, while several covariates display heterogeneous or even opposite effects across local contexts.

1. Introduction

Voting concentration refers to the degree to which electoral support is focused on a limited number of parties or candidates, reflecting how strongly voters' preferences are concentrated within the political system. High voting concentration occurs when a small number of parties or candidates receive the vast majority of votes, whereas voting fragmentation describes a more dispersed

[☆] This article is part of a Special issue entitled: 'SPASTA_Electoral Studies' published in Spatial Statistics.

* Corresponding author.

E-mail address: raffaele.mattera@unicampania.it (R. Mattera).

allocation of votes across multiple political actors. Understanding the drivers of voting concentration is therefore crucial for explaining political competition, electoral representation, and the structure of party systems.

In electoral studies, the concept of voting concentration is closely related to the literature on party-system fragmentation and the effective number of parties. A highly concentrated electoral outcome typically corresponds to a situation in which electoral support is captured by a limited number of parties or candidates, whereas fragmented systems are characterized by a more dispersed allocation of votes across multiple political actors. In this sense, voting concentration can be interpreted as the inverse counterpart of electoral fragmentation. A large body of literature has developed measures aimed at capturing this dimension, including the effective number of parties proposed by Laakso and Taagepera (1979), which measures party-system fragmentation by weighting parties according to their vote shares. Following this, Taagepera (1999) introduced complementary indicators designed to better capture asymmetries in party systems and the relative dominance of major parties, highlighting some limitations of relying exclusively on the classical effective number of parties index. More recently, Golosov (2010) proposed an alternative formulation that assigns greater weight to dominant parties while reducing the influence of very small parties, thereby improving the interpretability of the index in highly fragmented electoral systems characterized by the presence of numerous marginal political actors. These measures are widely used to summarize the degree of concentration or dispersion in electoral competition and to evaluate the structure of party systems.

Recent contributions emphasize that electoral fragmentation may have important behavioral and institutional consequences. For example, Zagórski (2022) shows that highly fragmented political environments may reduce electoral turnout by increasing informational complexity and voter uncertainty. Moreover, Krashinsky (2024) shows that local concentrations of voters may significantly affect voting participation, in particular participation rates increase where there exist greater concentrations of voters supporting majority parties. Moreover, Vitale (2025) shows that areas with a higher concentration of votes for a single council candidate receive more public goods. Other research investigates the causes of voting concentration. A large body of previous studies attributes voting concentration primarily to the party system fragmentation (Best, 2013; Meka, 2022; Valentim and Dinas, 2024). Among the institutional determinants of party system fragmentation, electoral systems are traditionally ascribed a central role. Specifically, electoral systems that enable or incentivize the presence of more parties tend to lead to higher fragmentation in voting outcomes. For example, Saarimaa and Tukiainen (2016) and Singer and Gershman (2018) show that changes in district magnitude can reduce the number of parties winning seats, while Golosov (2017) shows that electoral rules that emphasize the personal vote tend to increase fragmentation. Although there is consensus that a higher number of parties tends to produce more fragmented voting patterns, other research suggests that, under specific institutional settings, such outcomes may also result from candidate-level strategic behavior (e.g. see Cheibub et al., 2022).

While this literature has focused largely on macro-level institutional configurations, relatively little is known about how electoral fragmentation varies in response to characteristics at the local level, such as the socioeconomic and demographic profiles of municipalities, or the personal attributes of candidates (e.g., age, education, incumbency). These micro-level and contextual dimensions remain underexplored, despite their potential relevance in shaping the structure of electoral competition. Moreover, a growing body of scholars warns that studies of electoral outcomes that overlook spatial dependence risk producing biased or incomplete conclusions (Lacombe and Shaughnessy, 2007; Jensen et al., 2013). Spatial externalities and contextual effects have been documented in various electoral geography studies, often through spatial econometric techniques (Shin and Agnew, 2016). For example, Reményi et al. (2023) shows how voting outcomes in one locality are shaped by the dynamics of neighboring areas.

Despite the increasing attention paid to spatial factors in political behavior, the spatial dynamics of voting concentration remain substantially underexplored. While closely related outcomes, such as voter turnout, have received substantial attention from a spatial perspective (Cerqua and Di Matteo, 2025; Mansley and Demšar, 2015; Maškarinec, 2024), this is not yet the case for voting concentration. This paper aims to contribute to filling this gap by investigating the determinants of voting concentration at the micro-level considering data on Italian municipalities. In particular, we investigate to what extent the phenomenon is spatially heterogeneous across Italian municipalities and whether voting concentration in one municipality influences that observed in neighboring ones. To this end, we leverage data sourced from the Italian Ministry of the Interior about the 2018 elections and, following previous research (Maranzano et al., 2025; Monturano et al., 2025), a set of detailed municipality-level features.

From a statistical perspective, our research question needs the combination of clustering techniques and spatial autoregressive modeling. Indeed, given that the clustering structure of the Italian municipalities in terms of determinants of voting concentration is unknown ex-ante, we need a proper statistical procedure that simultaneously addresses both tasks, that is clustering units and estimating the spatial spillover effect. In this regard, Cerqueti et al. (2025b) recently proposed a novel statistical method, denoted as the Spatially-clustered Spatial Autoregressive (SCSAR) model, which solves both these issues iteratively. Specifically, the algorithm estimates spillover effects within each cluster – conditioning on some covariates – by means of SAR model (Elhorst et al., 2014) and estimates the group membership by maximizing a spatially-penalized likelihood function, thus extending the spatially-clustered regression (Sugasawa and Murakami, 2021) to spatial autoregression modeling. Our contribution is threefold:

- First, we study the local (i.e., group-specific) determinants of voting concentration employing entropy-based measures recently adopted in the literature (Marmani et al., 2020) to capture the complexity of party systems at the municipal level;
- Second, we explicitly model the spatial nature of voting concentration by the SCSAR model (Cerqueti et al., 2025b) and therefore provide the first empirical estimation of spatial spillovers in electoral fragmentation across Italian municipalities. In doing so, we also identify for the first time clusters of Italian municipalities characterized by the same degree of spillover in the voting concentration patterns.
- Third, we make the SCSAR framework operational for a large municipal dataset by exploiting sparse matrix factorization and by documenting, through a dedicated computational benchmark, how the proposed implementation scales relative to the standard eigenvalue-based approach.

The implementation of the spatially-clustered version of the SAR is grounded in both methodological, empirical and computational reasons. A standard SAR model is appropriate when the objective is to estimate a single average spatial spillover effect under the assumption that both the autoregressive parameter and the regression coefficients are homogeneous across all spatial units. In our setting, however, this assumption might be too restrictive. Therefore, in line with the spatial heterogeneity framework (Goodchild, 2004; Zhu and Turner, 2022), our substantive objective is not only to test whether voting concentration displays spatial dependence, but also to assess whether such dependence and the associated covariate effects vary across groups of municipalities characterized by different local political, demographic, and socio-economic conditions. This is especially relevant in the Italian case, where territorial heterogeneity is a central empirical feature of the phenomenon under study (Maranzano et al., 2025). In this regard, unlike a global SAR specification, the SCSAR specification allows for cluster-specific spatial autoregressive coefficients and cluster-specific regression parameters, thus capturing spatial heterogeneity in both spillover intensity and local determinants. In other words, it relaxes the implicit assumption that neighboring effects and covariate-response relationships are constant throughout the country. Moreover, the spatially penalized clustering step promotes spatially coherent groups, which is important for substantive interpretation; in particular, estimating a cluster-specific spillover effect over highly fragmented and geographically disconnected sets of municipalities would be much harder to interpret as a meaningful local spatial regime. In this sense, the clustering component is not introduced only to improve fit, but also to ensure that the estimated local spillovers remain geographically interpretable.

However, we remark that the SCSAR algorithm as proposed in Cerqueti et al. (2025b) might struggle in being estimated with very large number of spatial units n (i.e., $n \gg 0$), as in the case of Italian municipalities. Indeed, the maximization of the log-likelihood function required in the estimation of spatial autoregressive models is computationally demanding, as it involves the evaluation of the log-determinant of an $n \times n$ matrix. The standard approach based on eigenvalue decomposition (Ord, 1975) is computationally intensive also for relatively small number n . In our setting, dealing with over $n = 8000$ Italian municipalities, the computational burden is particularly important. To ensure scalability, we adopt exact methods based on Cholesky factorization (Pace, 1997; Pace and Barry, 1997b), exploiting the sparsity of the spatial weights matrix. Although the usefulness of factorization techniques is well established in spatial statistics, this point is not entirely trivial in the clustered setting. On the one hand, clustering decomposes the global spatial system into K smaller subproblems, which can improve scalability when clusters are not excessively unbalanced. On the other hand, the algorithm repeatedly rebuilds cluster-specific spatial weights, estimates multiple SAR models, and updates labels across iterations. Therefore, we also provide a computational benchmark showing the usefulness of factorization techniques in large n setting for this class of models. This is useful not to re-establish the general superiority of sparse factorization, but to document how these computational gains materialize within the complete clustered SAR estimation procedure.

From the computational perspective, our results show that sparse factorization consistently improves the scalability of SCSAR estimation and makes the model empirically feasible for large spatial systems such as the Italian municipal graph. The computational study confirms that the sparse implementation substantially reduces execution time relative to the standard eigenvalue-based approach, especially as the number of spatial units increases. From the empirical perspective, we find that electoral concentration in Italy is shaped by a combination of structural factors, such as income, demographic composition, and population density, that prove robust across models, while other determinants vary depending on the territorial context. Spatial spillovers remain consistently positive, although the global estimates are unreasonably large compared to those obtained with clustered models. By explicitly modeling spatial dependence, we demonstrate that electoral behavior cannot be understood solely through municipal characteristics, as strong and persistent spillovers reveal the territorial contagion of voting patterns.

The rest of the paper is structured as follows. Section 2 describes the data adopted in our paper and explains in detail the statistical measures adopted for measuring voting concentration patterns. Section 3 presents the spatially-clustered spatial autoregressive model adopted in the paper and the scalable algorithm used to estimate it, while Section 5 discusses the results obtained. Section 6 concludes the paper with final remarks and future research directions.

2. Electoral Italian data

Our empirical analysis is based on one of the most comprehensive databases of Italian elections, provided by the Italian Ministry of the Interior.¹ The raw dataset was transformed into a structured SQL database, divided by election type: Municipal, Provincial, Regional, Chamber of Deputies, Senate, and European Parliament. Second-round runoffs (introduced by Law No. 81 of 1993) in municipal and provincial elections were excluded from the analysis. The data has been further organized into two subsets for each election: (i) party and candidate information, and (ii) electoral outcomes. To enable spatial analysis, the data were restructured using municipalities as the primary geographic unit. This fine spatial scale reduces the risk of masking local heterogeneity through aggregation, a well-known issue in spatial analysis commonly discussed as the modifiable areal unit problem (MAUP), whereby spatial patterns and statistical associations may depend on the chosen areal unit (Arbia and Petrarca, 2011; Manley, 2021).

Our analysis focuses on the Italian Chamber of Deputies and Senate, which together constitute the Italian bicameral legislative system. The 2018 general election was held under Law no. 165/2017, commonly referred to as the *Rosatellum*, which introduced a mixed-member electoral system combining proportional representation and plurality rule. In both chambers, a share of seats was allocated through single-member constituencies under a winner-take-all mechanism (i.e., *majoritarian* mechanism), while the remaining seats were distributed proportionally within multi-member constituencies (i.e., *proportional* mechanism). The electoral system also included electoral thresholds and did not allow split-ticket voting (Camera dei Deputati & Senato della Repubblica

¹ We refer the reader to the link <https://elezionistorico.interno.gov.it> (Accessed on September 24th, 2025) to access the online open source database.

Italiana, 2022). Despite these common institutional features, the two chambers differed in several politically relevant dimensions. First, the territorial organization of the electoral system was not identical: the Senate was divided into 20 electoral districts corresponding to the Italian administrative regions, whereas the Chamber of Deputies was organized into 28 districts, as the four largest regions were further subdivided. Each district was then partitioned into single-member and multi-member constituencies. Second, until 2021 the electorate differed across chambers: all citizens aged 18 and above were eligible to vote for the Chamber of Deputies, while only citizens aged 25 and above could vote for the Senate. As a consequence, the Senate electorate represented an older and more restricted subset of the Chamber electorate. Third, although the formal electoral thresholds were broadly similar across the two chambers (i.e., a 3% threshold for party lists and a 10% threshold for coalitions), Senate seats were constitutionally allocated on a regional basis rather than at the national level. Moreover, the Senate was characterized by a smaller total number of seats and, consequently, by lower effective district magnitudes within proportional constituencies. These characteristics likely implied higher effective electoral thresholds in several regions and are generally associated with a weaker correspondence between vote shares and seat allocation (i.e., *disproportionality*), as well as with stronger incentives for strategic coordination among voters and parties, potentially reducing the electoral viability of smaller parties. Overall, these institutional and territorial differences may contribute to explaining territorial heterogeneity in voting concentration patterns and may partially account for the differences observed in the empirical results across the two chambers.

Moreover, recent research has shown that the COVID-19 pandemic profoundly altered electoral behavior, affecting both turnout and party preferences. Empirical evidence from Spain and France indicates that local exposure to COVID-19 reduced participation and shifted support towards nationalist or pro-/anti-lockdown parties, while stricter lockdown policies mobilized voters and increased support for incumbents (Fernandez-Navia et al., 2021; Lall and Vilalta, 2025; Giommoni and Loumeau, 2022). To avoid conflating such crisis-induced dynamics with longer-term political and territorial patterns, our analysis deliberately focuses on the last Italian general election held before the pandemic, that is, in the year 2018.

2.1. Voting concentration measures: definitions and literature review

To assess electoral concentration at the municipal level, we rely on two complementary distributional indicators, namely the Shannon entropy index and the Gini coefficient.

Entropy measures have been widely used in economics, finance, sociology, and information theory to quantify disorder, uncertainty, or heterogeneity in systems (Fabretti and Ausloos, 2005; Cerqueti et al., 2018; Lambiotte and Ausloos, 2006; Dragan and Isaac-Maniu, 2019). In electoral studies, the Shannon index has been applied to model political uncertainty and voter behavior (Hooghe and Teepe, 2007; Gill, 2005; Borghesi et al., 2012; Coleman, 2004), though often at the national level or with survey data.

More formally, we consider n municipalities (with $i = 1, \dots, n$). The entropy for municipality i in a given election is defined as

$$H_i = - \sum_{j=1}^{M_i} p_j \log_{M_i} p_j, \quad \forall i = 1, \dots, n \quad (1)$$

where $p_j > 0$ is the vote share of party j , and M_i is the number of candidate parties in that municipality that collected at least one vote in municipality i . This formulation naturally normalizes the standard Shannon entropy by the number of competing parties, so that $0 \leq H_i \leq 1$. Recall that, in its standard definition, the entropy increases when votes are evenly distributed across parties, and decreases when one party tends to dominate over the others strongly. In contrast, the Gini coefficient captures inequality in the distribution of votes, increasing when one or a few parties concentrate most of the support. In order to compare them directly, we adopt a polarized transformation of the entropy, that is, $H'_i = 1 - H_i$, such that higher values indicate higher concentration. This simple polarization aligns its interpretation with that of the Gini coefficient.


The Gini coefficient, originally developed to measure income inequality (Gini, 1921), has also been applied to political science, electoral clustering, and machine learning (Cerqueti and Ausloos, 2015; Manek et al., 2017; Walks et al., 2014). While entropy measures general dispersion, the Gini coefficient captures skewness or polarization, thus providing a complementary view of electoral inequality. Its combination with entropy-based measures has been proposed to jointly assess polarization and complexity in voting systems (Preda et al., 2015; Tanak et al., 2017). The Gini index computed for municipality i referring to a generic election is given by the formula

$$G_i = \frac{\sum_{j=1}^{M_i} \sum_{h=1}^{M_i} |x_j - x_h|}{2M_i \sum_{j=1}^{M_i} x_j}, \quad (2)$$

where x_j and x_h are vote counts for parties j and h in the generic municipality i . Since this index may be biased in small samples or elections with few competitors (Deltas, 2003; Raffinetti et al., 2015), we normalize it with respect to its maximum attainable value under the extreme scenario where one party receives all votes and the remaining $M_i - 1$ parties receive zero votes. In this latter case the maximum attainable Gini is $U_i = \frac{M_i - 1}{M_i}$ and the raw Gini coefficient G_i is bounded between $0 \leq G_i \leq \frac{M_i - 1}{M_i}$. Therefore,


in the following, we adopt a normalized version of Gini index, defined as $G_i^{adj} = \frac{G_i}{U_i}$, that bounds the range of the index between $0 \leq G_i^{adj} \leq 1$ and allows for complete comparability of voting inequality across municipalities with different numbers of candidates.

It is worth of mention that the two concentration measures grounds on different mathematical and statistical properties that could generate substantial differences in the empirical results. Eq. (2) clearly shows that the normalized Gini index is primarily driven by pairwise differences among vote shares and therefore mainly reflects the extent to which the electoral distribution is



(a) Gini index for Deputies Chamber

(b) Gini index for Senate Chamber

Fig. 1. Municipal-level Gini indices for the two Chambers.

(a) Polarized and normalized Shannon-based concentration index) for Chamber of Deputies

(b) Polarized and normalized Shannon-based concentration index for the Senate

Fig. 2. Municipal-level Shannon-based concentration indices for the two Chambers.

dominated by one or a few large parties. Consequently, the index is disproportionately influenced by major political forces whose electoral support is typically organized around persistent territorial political structures, long-standing ideological traditions, stable party organizations, and regionally rooted voting cultures. These mechanisms tend to evolve smoothly over space and are therefore more likely to generate strong spatial dependence patterns. By contrast, the Shannon index defined in Eq. (1) is an entropy-based measure that captures the overall dispersion of vote shares across the entire party system. Unlike the Gini index, the Shannon measure remains highly sensitive not only to dominant parties, but also to variations in the relative importance of small and medium-sized parties. Indeed, based on the above notation, the marginal sensitivity of the Shannon entropy with respect to the i th party within municipality j can be defined as $\partial H_i / \partial p_j = -\frac{1 + \log p_j}{\log(M_j)}$. In practical terms, this latter expression implies that even modest fluctuations in the electoral performance of minor lists, local candidates, regional movements, or municipality-specific political configurations (i.e., when $p_j \rightarrow 0$) may induce comparatively large reactions and substantially affect the index. Such fine-grained fragmentation dynamics are generally more localized, context-dependent, and spatially unstable across neighboring municipalities. Consequently, Shannon-type measures – incorporating a larger idiosyncratic local component – might amplify municipality-specific fragmentation generated by minor parties and local electoral supply (therefore exhibiting weaker spatial autoregressive structures), whereas Gini-type measures mostly reflect broader and spatially persistent territorial political structures.

Considering our data related to 2018 elections, both indicators reveal systematic territorial divides. Figs. 1 and 2 show the spatial heterogeneity of voting concentration.

The electoral concentration tends to be lower in southern municipalities, reflecting greater party fragmentation, yet heterogeneity also emerges within macro-areas, with differentiated patterns in both North and South. These findings justify the need for clustering strategies that explicitly account for spatial heterogeneity (Goodchild, 2004; Zhu and Turner, 2022). Therefore, the Italian case is

Table 1
Variable names, acronyms, and brief descriptions.

Variable	Acronym	Description
Panel A: dependent variables		
Gini (Chamber of Deputies)	Gini_c	Inequality of votes in Chamber elections (Gini index)
Shannon (Chamber of Deputies)	Shannon_c	Entropy of vote distribution in Chamber elections
Gini (Senate)	Gini_s	Inequality of votes in Senate elections (Gini index)
Shannon (Senate)	Shannon_s	Entropy of vote distribution in Senate elections
Panel B: covariates by type		
<i>Economic:</i>		
Average income	AvgInc	Mean income per capita in the municipality
Gini (Income)	Gini	Gini coefficient for income per capita
Employment rate (20–64)	ER	Share of employed population aged 20–64
Land use (ventile)	LUind	Land use density for industrial/service areas (ventile ranking)
% working in low-productivity sectors	WLP	Share of workers in low productivity sectors (ventile ranking)
Share of Poverty	SoP	Share of population declaring less than 10 thousand € in 2018
<i>Demographic:</i>		
Population density	POPd	Residents per square kilometer
Population Dependency Index	PDI	Ratio of dependent age groups to working-age population
Foreigner Population	FPop	Share of foreigners of total population
<i>Institutional:</i>		
Mayor has degree	Degree	Binary indicator for whether mayor has a university degree
Mayor age	Age	Age of the mayor at the time of election
Access to essential services index	ESI	Composite index on access to basic public services

particularly suited for this type of analysis for at least two reasons. First, Italy exhibits significant territorial heterogeneity in either economy, society, and political patterns (Faggian et al., 2021; Maranzano et al., 2025). Regional disparities between the North and South (Cerqueti et al., 2025a; Musolino, 2018), urban–rural divides (Ballatore and Mariani, 2019; Bazzoli and Lello, 2022), and the distinct historical and cultural identities of various macro-areas contribute to differentiated political behaviors and electoral outcomes across space. This heterogeneity suggests that local determinants of concentration are unlikely to be uniform throughout the country. Second, the Italian party system has undergone frequent and radical transformations over the last decades (Shin and Agnew, 2007), including party splits, mergers, the rise of new political forces, and electoral reforms. These changes have often had uneven territorial impacts, with some areas becoming more politically volatile or fragmented than others (Marmani et al., 2020). Hence, it is reasonable to assume that different groups of municipalities exist, and that these groups exhibit distinct configurations of determinants, reinforcing the need for spatially sensitive empirical strategies.

2.2. Additional municipal-level data

To analyze the determinants of voting concentration at the municipal level, it is essential to consider a range of economic, demographic, and institutional characteristics that may shape local electoral dynamics. These variables capture structural conditions of the territory and political context that can influence both the political supply and voter preferences. Table 1 briefly describes the selected variables for each of these dimensions (for a comprehensive description of the available data we refer the reader to the paper by Maranzano et al., 2025, which also provides an extended dataset in the supplementary materials. In particular, we stress that their spatio-temporal dataset was assembled joining municipal-level data from multiple Italian statistical sources, such as Ministry of Economy and Finance, the Ministry of the Interior and the National Statistics Office (ISTAT).).

Economic variables. The economic structure of a municipality can affect both the credibility of the local government and the preferences of voters (Truglia and Zeli, 2020). Moreover, structural fragility amplifies the demand for outsider or protest platforms (Guriev and Papaioannou, 2022). In this spirit, we select the following variables. The average income per capita (AvgInc) acts as a barometer of local prosperity and fiscal capacity. The share of low-income taxpayers (SoP) captures deprivation on the extensive margin; where SoP is high, we expect protest voting to be more common, and fragmentation to rise. Income inequality (IncGini) introduces distributional tensions and a wider set of competing claims. The employment rate among those aged 20–64 (ER) summarizes labor-market health. Finally, we consider production through two rank-based indicators: the national ventile of industrial and service land use (LUind), a proxy for functional and economic density that typically favors consolidation and concentration, and the national ventile of employment in low-productivity sectors (WLP), a marker of structural fragility commonly associated with outsider or protest vote and thus with fragmentation.

Demographic variables. Demography and settlement patterns shape the heterogeneity of preferences (Gunther et al., 2015). For example, social capital and district magnitude well predict the share of variation in preference voting (Passarelli, 2017). Therefore, we consider population density (POPd) and the population dependency index (PDI). Moreover, we consider the share of foreign residents (FPop) which reflects socio-demographic diversity and the salience of migration-related issues (Barone et al., 2016).

Institutional variables. As the link between voters and parties has weakened, voters have substituted their partisanship with other factors in their electoral decision-making process, notably the assessment of party candidates' characteristics (Coffé and Von Schoultz, 2021). To account for such an aspect, we consider the following variables. First, we collect information on the mayor's

Table 2

Main descriptive statistics of the variables associated with this study. See Table 1 for explanation of the acronyms.

Variable	Mean	St. Dev.	Min.	Max.
Panel A: dependent variables				
Gini_c	0.69	0.07	0.11	0.86
Gini_s	0.69	0.07	0.20	0.86
Shannon_c	0.34	0.06	0.02	0.82
Shannon_s	0.35	0.06	0.05	0.78
Panel B: covariates				
AvgInc	17 511.56	3740.84	5800.18	45 645.12
Gini	0.31	0.04	0.02	0.46
FPop	6.57	4.17	0.00	38.75
SoP	31.43	10.08	11.11	77.98
POPd	2.24	5.27	0.01	130.45
Degree	0.54	0.50	0.00	1.00
Age	58.86	9.58	27.00	93.00
ESI	30.93	18.12	0.00	317.60
PDI	72.70	9.94	28.00	187.50
ER	64.54	9.68	36.09	91.26
LUind	10.47	5.76	1.00	20.00
WLP	10.57	5.75	1.00	20.00

education. If the mayor holds a university degree (Degree), it provides a compact proxy for candidate quality and administrative human capital. Moreover, mayoral age (Age) may signal experience or generational distance, as younger leaders can mobilize new cohorts while older leaders may consolidate traditional bases. Finally, we consider the index of access to essential services (ESI), which summarizes the effectiveness of local provision and infrastructure shaping voting behaviors.

Table 2 shows the main descriptive statistics associated with the selected covariates.

3. Methodology

This section introduces the methodological framework employed throughout the paper. Section 3.1 reviews the classical spatial autoregressive regression model. Section 3.2 presents the proposed regression framework based on spatial clustering, while Section 3.4 describes the computational strategy adopted to efficiently estimate the SCSAR model in settings characterized by a large number of spatial units.

3.1. The spatial autoregressive (SAR) model

Let us consider n spatial units, in our case the Italian municipalities, for which we observe a dependent variable $\mathbf{y} = [y_1, y_2, \dots, y_i, \dots, y_n]'$ and a set of p covariates $\mathbf{x}_p = [x_{p1}, x_{p2}, \dots, x_{pi}, \dots, x_{pn}]'$, collected in a matrix $\mathbf{X}$ of dimension $(n \times p)$. The dependent variable is one of the voting concentration measures described in Section 2.1, while the set of covariates is that presented in Section 2.2. Let us start with the linear model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n), \quad (3)$$

where $\mathbf{y} \in \mathbb{R}^n$, $\mathbf{X} \in \mathbb{R}^{n \times p}$, and $\boldsymbol{\beta} \in \mathbb{R}^p$, with $\boldsymbol{\beta}$ the set of regression coefficients and $\boldsymbol{\varepsilon}$ the error term. Spatial dependence can be introduced by modifying (3) by means of the following (spatial) autoregressive specification

$$\mathbf{y} = \rho \mathbf{W}\mathbf{y} + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (4)$$

with scalar spatial autoregressive parameter ρ and a given sparse spatial weights matrix $\mathbf{W} = (w_{ij})_{i,j} \in \mathbb{R}^{n \times n}$ with zero diagonal elements. The spatial weight matrix reflects the neighboring structure of the n spatial units, and it is known. Specifically, we assume that the elements w_{ij} are binary, that is,

$$w_{ij} = \begin{cases} 1 & \text{if unit } i \text{ is contiguous to unit } j, \\ 0 & \text{otherwise.} \end{cases}$$

Throughout, we employ a subsequently row-standardized matrix $\mathbf{W}$. By moving the spatial-lag term $\rho \mathbf{W}\mathbf{y}$ to the left-hand side, we get $\mathbf{y} - \rho \mathbf{W}\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$. Let us define $\mathbf{A}(\rho) = \mathbf{I}_n - \rho \mathbf{W}$, then

$$\mathbf{A}(\rho)\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}. \quad (5)$$

Under regularity conditions, the parameter space for ρ ensures non-singularity of $\mathbf{A}(\rho)$, with typical bounds involving the inverse of the spectral radius of $\mathbf{W}$. The log-likelihood associated with the SAR model, $\ell(\rho, \boldsymbol{\beta}, \sigma^2)$, is given by (Elhorst et al., 2014)

$$\ell(\rho, \boldsymbol{\beta}, \sigma^2) = \log |\mathbf{A}(\rho)| - \frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} (\mathbf{y} - \rho \mathbf{W}\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' (\mathbf{y} - \rho \mathbf{W}\mathbf{y} - \mathbf{X}\boldsymbol{\beta}). \quad (6)$$

Alternatively, by writing the i th component of vector $\mathbf{r}(\rho, \beta) = \mathbf{y} - \rho \mathbf{W} \mathbf{y} - \mathbf{X} \beta$ as $r_i(\rho, \beta) = y_i - \rho \sum_j w_{ij} y_j - \mathbf{x}_i' \beta$ for each $i = 1, \dots, n$, we can write the likelihood as follows:

$$\ell(\rho, \beta, \sigma^2) = \log |\mathbf{A}(\rho)| - \frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \sum_{i=1}^n r_i(\rho, \beta)^2. \quad (7)$$

The model parameters can be estimated with maximum likelihood. The main issue with the SAR (4) is that a single global model may be mis-specified when spatial relationships between $\mathbf{y}$ and $\mathbf{X}$ are heterogeneous and when the spatial autocorrelation is not constant across the n units. Cerqueti et al. (2025b) provides a solution to estimate heterogeneous SAR models at the cluster level.

3.2. The clusterwise approach: SCSAR

Let the n spatial units be partitioned into K clusters ($k = 1, \dots, K$), with $\mathbf{k} = (k_1, \dots, k_n)$ the vector of labels (i.e., $k_j \in \{1, \dots, K\}$) is the cluster assignment of unit j , for each $j = 1, \dots, n$. For cluster k , denote the index set by $I_k = \{i : k_i = k\}$ with size $n_k = |I_k|$, and let $(\mathbf{y}_k, \mathbf{X}_k)$ be the restriction of $(\mathbf{y}, \mathbf{X})$ to I_k . Write $\mathbf{W}_{kk}$ for the $(n_k \times n_k)$ submatrix of $\mathbf{W}$ induced by I_k , and define the quantity $\mathbf{A}_k(\rho_k) = \mathbf{I}_{n_k} - \rho_k \mathbf{W}_{kk}$. We clarify that $\mathbf{W}_{kk}$ is not obtained by simply taking the corresponding block of a globally row-standardized matrix. Rather, for each cluster k , we first extract the binary subgraph induced by I_k from the original adjacency matrix and then construct a cluster-specific row-standardized matrix.

Given the clustering structure $\mathbf{k}$, the k th clustered SAR model is described as

$$\mathbf{A}_k(\rho_k) \mathbf{y}_k = \mathbf{X}_k \beta_k + \varepsilon_k, \quad \varepsilon_k \sim \mathcal{N}(\mathbf{0}, \sigma_k^2 \mathbf{I}_{n_k}), \quad \forall k = 1, \dots, K \quad (8)$$

having log-likelihood

$$\ell_k(\rho_k, \beta_k, \sigma_k^2) = \log |\mathbf{A}_k(\rho_k)| - \frac{n_k}{2} \log(2\pi\sigma_k^2) - \frac{1}{2\sigma_k^2} \sum_{i=1}^{n_k} r_{ik}(\rho_k, \beta_k)^2, \quad (9)$$

where

$$r_{ik}(\rho_k, \beta_k) = y_i - \rho_k \sum_j [\mathbf{W}_{kk}]_{ij} y_j - \mathbf{x}_i' \beta_k.$$

with $[\mathbf{W}_{kk}]_{ij}$ being the spatial weight associated with units i and j within the k th cluster. The cluster-specific parameters $\delta_k = (\rho_k, \beta_k, \sigma_k^2)$ are estimated by cluster-wise maximum likelihood. This conditional row-standardization is used only in the estimation of the cluster-specific SAR models. Consequently, ρ_k measures the strength of spatial dependence within the interaction domain induced by cluster k . It should therefore not be interpreted as numerically identical to the global SAR coefficient, which is defined on the full spatial graph. However, because each $\mathbf{W}_{kk}$ is normalized within the corresponding cluster, the resulting ρ_k values retain a meaningful relative interpretation across clusters.

Because the clustering partition $\mathbf{k}$ is unknown ex-ante, Cerqueti et al. (2025b) proposed to estimate jointly the labels and the parameters, following a similar rationale of the spatially clustered regression model (SCR, Sugawara and Murakami, 2021). Let us consider the following penalized likelihood

$$Q(\delta, \mathbf{k}) = \sum_{k=1}^K \ell_k(\delta_k) + \phi \sum_{i < j} w_{ij} \mathbb{I}(k_i = k_j), \quad (10)$$

where $\delta = (\delta_1, \dots, \delta_K)$ is the set of parameters in the K clusters and $\phi \geq 0$ controls spatial smoothness of the labels. The SCSAR algorithm alternates the following two steps until convergence of Q or stabilization of the labels:

1. (*Estimation step*). Given a provisional partition $\mathbf{k}^{(t)}$, for each cluster k maximize $\ell_k(\delta_k)$ with respect to δ_k on each of the K clusters, obtaining $\hat{\delta}_k^{(t)}$ for each $k = 1, \dots, K$;
2. (*Clustering step*). Holding $\hat{\delta}_k^{(t)}$ fixed, reassign each unit i to the cluster that yields the largest increase in Q .

For computational efficiency, Cerqueti et al. (2025b) use a per-unit score that decomposes the cluster likelihood additively, that is,

$$\ell_{ik}(\delta_k) = \frac{1}{n_k} \log |\mathbf{A}_k(\rho_k)| - \frac{1}{2} \log(2\pi\sigma_k^2) - \frac{1}{2\sigma_k^2} r_{ik}(\rho_k, \beta_k)^2, \quad (11)$$

such that $\sum_{i \in I_k} \ell_{ik}(\delta_k) = \ell_k(\delta_k)$. The proposed label for unit i is

$$k_i = \arg \max_{k \in \{1, \dots, K\}} \left\{ \ell_{ik}(\hat{\delta}_k^{(t)}) + \phi \sum_j w_{ij} \mathbb{I}(k_j^{(t)} = k) \right\}. \quad (12)$$

It is also worth noting that the matrix used in the SAR likelihood and the graph used in the spatial penalty play different roles. The likelihood contribution $\ell_k(\delta_k)$ is computed using the cluster-specific row-standardized matrix $\mathbf{W}_{kk}$, whereas the penalty term in (10) is computed on the original global adjacency graph. The latter is used only to encourage neighboring municipalities to share the same label and is not re-standardized within clusters.

We also remark that the SAR likelihood is joint and does not really factorize by units. Nevertheless, (11) provides an additive decomposition of ℓ_k that assigns to each unit a share of the Jacobian and the normalizing constant (both common within the

cluster) plus its own quadratic residual contribution. Therefore, the loss to be optimized is a computationally feasible surrogate of (10). This construction is crucial for scalable label updates; indeed, it yields a local score whose sum exactly matches the cluster likelihood, so that comparisons across candidate clusters for unit i are aligned with the global objective. Practically, the SCSAR algorithm uses (12) to evaluate label changes, and assign the i th unit to the cluster allowing maximizing the surrogate of the total penalized likelihood (10). Therefore, this approach preserves the correctness of maximum-likelihood optimization while keeping the clustering step computationally tractable. The procedure is shown in Algorithm 1. At each iteration, given the current estimates $\hat{\delta}_k^{(t)}$ and the labels of the neighboring units, the criterion in (12) is separable across units, so that the assignment step attains its exact maximum by pointwise comparison: unit i changes label only if this increases its own penalized score. Algorithm 1 therefore belongs to the family of alternating optimization procedures, and by construction it converges to a local optimum of the surrogate objective. Global optimality is not guaranteed, which is why multi-start initialization is adopted. The first initialization is based on a standard k -means procedure involving all variables and the spatial coordinates. This choice provides a simple starting rule that combines feature similarity and geographical proximity. To reduce sensitivity to a single random start of the k -means, the procedure is repeated multiple times and the partition associated with the smallest within-cluster sum of squares is retained as the initial allocation for the iterative SCSAR updates. We remark that, as in main statistical software concerning spatial regression (e.g. see Bivand and Piras, 2015), for the numerical optimization ρ_k is constrained to a closed interval contained in $(-1, 1)$ and the concentrated log-likelihood is then maximized by constrained one-dimensional search over this interval. Moreover, units left without neighbors within their own cluster are handled by setting their spatial lag to zero, although the spatial penalty makes such configurations infrequent.

Algorithm 1 Fast SCSAR estimation

Require: Binary global adjacency matrix $\mathbf{W}$; number of clusters K ; penalty parameter $\phi \geq 0$; maximum number of iterations $T_{\max}$; tolerances $\text{RelTol}, \text{AbsTol}$.

Ensure: Final partition $\mathbf{k}$, cluster-specific estimates $\hat{\delta}_1, \dots, \hat{\delta}_K$ and the loss Q .

1: Initialize $\mathbf{k}^{(0)}$ and set $t \leftarrow 0$, $Q^{(0)} \leftarrow -\infty$.

2: **while** $t < T_{\max}$ **do**

3: $t \leftarrow t + 1$.

4: **for** $k = 1, \dots, K$ **do**

5: $I_k^{(t-1)} = \{i : k_i^{(t-1)} = k\}$.

6: Build the row-standardized cluster-specific weights matrix $\mathbf{W}_{kk}^{(t-1)}$ on the subgraph induced by $I_k^{(t-1)}$.

7: Estimate

$$\hat{\delta}_k^{(t)} = \arg \max_{\delta_k} \ell_k(\delta_k; \mathbf{y}_k, \mathbf{X}_k, \mathbf{W}_{kk}^{(t-1)}),$$

8: **end for**

9: **for** $i = 1, \dots, n$ **do**

10: **for** $k = 1, \dots, K$ **do**

11: Define the candidate row-standardized weights

$$\tilde{w}_{ij,k}^{(t-1)} = \frac{w_{ij} \mathbb{I}(k_j^{(t-1)} = k)}{\sum_{\ell=1}^n w_{i\ell} \mathbb{I}(k_\ell^{(t-1)} = k)},$$

12: Compute the per-unit log-likelihood contribution $\tilde{\ell}_{ik}^{(t)}$ as in (11), evaluated at the current iteration estimates $\hat{\delta}_k^{(t)}$ and at the candidate weights at the previous iteration $\tilde{w}_{ij,k}^{(t-1)}$. Then, the penalized score is

$$Q_{ik}^{(t)} = \tilde{\ell}_{ik}^{(t)} + \phi \sum_{j=1}^n w_{ij} \mathbb{I}(k_j^{(t-1)} = k).$$

13: **end for**

14: Update the partition

$$k_i^{(t)} = \arg \max_{k \in \{1, \dots, K\}} Q_{ik}^{(t)}.$$

15: **end for**

16: Set

$$Q^{(t)} = \sum_{i=1}^n Q_{ik_i^{(t)}}^{(t)}.$$

17: **if** $t > 1$ and $\left| \frac{Q^{(t-1)} - Q^{(t)}}{Q^{(t)}} \right| \leq \text{RelTol}$ **or** $|Q^{(t)} - Q^{(t-1)}| \leq \text{AbsTol}$ **then**

18: **break**

19: **end if**

20: **end while**

21: **return** $\mathbf{k}$, $\hat{\delta}_1, \dots, \hat{\delta}_K$, and Q .

We highlight that a purely likelihood-driven reassignment, obtained by setting $\phi = 0$, may group together units that are far apart in space when their regression fits look similar. This weakens geographical interpretability, risks omitting non-negligible cross-border interactions in the within-cluster SAR, and degrades sparse computations by fragmenting $\mathbf{W}_{kk}$. The Potts-type penalty in (10), instead, rewards neighboring units sharing the same label, encouraging contiguous groups, reducing the number of edges cut by the partition, and aligning clusters with the underlying spatial graph. The spatial penalty parameter ϕ governs the trade-off between fit and spatial coherence: large values of ϕ push the solution towards compact and nearly contiguous clusters, whereas small

values of ϕ allow the likelihood component to dominate when the data support sharper or less spatially compact groupings. This soft-regularization perspective also clarifies the relationship between SCSAR and the broader literature on endogenous spatial regimes (Anselin and Amaral, 2024). In particular, SCSAR is related to the SKATER-regression approach discussed by Anselin and Amaral (2024), since both methods use model fit to endogenously construct spatial groups. However, they differ in the way geographical coherence is imposed. In SKATER-regression, contiguity is a hard constraint, that is, regimes are obtained by pruning a minimum spanning tree, which guarantees that the resulting groups consist of connected observations. This is a coherent modeling choice when spatial regimes are required to be contiguous by construction. By contrast, SCSAR adopts a less restrictive strategy: spatial coherence is encouraged through the Potts-type penalty entering the penalized likelihood objective, but it is not imposed as an absolute graph-theoretic constraint, making the degree of spatial compactness data-adaptive. In this sense, SCSAR provides a flexible compromise between regression fit, spatial dependence and geographical coherence.

Some final remarks on the theoretical status of the procedure are in order. For a fixed oracle partition, each cluster-specific specification reduces to a standard SAR model on the corresponding induced subgraph. Under the regularity conditions of Lee (2004), the associated within-cluster estimator therefore inherits the usual consistency and asymptotic normality results. As in related approaches to endogenous spatial regimes (Anselin and Amaral, 2024) and spatially clustered regression (Sugasawa and Murakami, 2021), however, correct recovery of the underlying regimes is assumed rather than established, while the calibration of (K, ϕ) relies on a heuristic rule. In this regard, we note that because the SCSAR partition is estimated from the same data, the reported standard errors are conditional on the selected partition and do not account for the additional uncertainty associated with regime estimation. They may therefore understate the overall uncertainty. Developing a full post-clustering inferential theory for SCSAR is beyond the scope of this work and is left for future research.

3.3. Hyperparameter selection: a heuristic procedure

The selection of (K, ϕ) is made by information criteria and therefore follows the same spirit of Sugawara and Murakami (2021) and Cerqueti et al. (2025b). A natural starting point would be to evaluate the whole penalized objective (10) at the estimated quantities, which yields the following *penalized* information criterion

$$\text{IC}^{\text{pen}}(K, \phi) = -2 \sum_{k=1}^K \ell_k(\hat{\delta}_k(K, \phi)) - 2\phi \sum_{i < j} w_{ij} \mathbb{I}(k_i = k_j) + p(K) \log n, \quad (13)$$

where $p(K) = K(p+2)$ counts the cluster-specific parameters $(\beta_k, \rho_k, \sigma_k^2)$ across the K clusters and n is the number of spatial locations. We write $\ell_k(\hat{\delta}_k(K, \phi))$ to make it clear that likelihood and the related estimates depend on the choice of the hyperparameters (K, ϕ) . The criterion (13), however, cannot be used to calibrate ϕ . Since $2\phi \sum_{i < j} w_{ij} \mathbb{I}(k_i = k_j) \geq 0$, the last term is non-positive and strictly decreasing in ϕ for any fixed partition. Moreover, larger values of ϕ induce more spatially compact partitions. Hence $\text{IC}^{\text{pen}}(K, \phi)$ decreases monotonically in ϕ through both channels, and its joint minimization over (K, ϕ) admits no interior solution: the selected value of ϕ is always the largest one available on the grid, irrespective of the data, thus returning an almost purely geographical partition.

For this reason, following previous works (Sugasawa and Murakami, 2021; Cerqueti et al., 2025b), we adopt the following standard non-penalized information criterion

$$\text{IC}(K, \phi) = -2 \sum_{k=1}^K \ell_k(\hat{\delta}_k(K, \phi)) + p(K) \log n, \quad (14)$$

in which ϕ does not appear explicitly and enters only implicitly, through the estimated partition $\hat{\mathbf{k}}(K, \phi)$ and hence through the cluster-specific estimates $\hat{\delta}_k(K, \phi)$. In this formulation ϕ is treated as a regularization parameter rather than as a free parameter to be optimized: since it does not add parameters to the model, a strictly positive ϕ is retained only when the spatially regularized partition attains a better fit at the same model complexity.

Operationally, (K, ϕ) are calibrated through a sequential elbow-based procedure. For each fixed value of ϕ on a suitably rescaled grid of small values proportional to the total number of spatial units,² we compute the elbow point of the corresponding curve (14) as a function of K , obtaining one candidate value K_ϕ^{elbow} for each ϕ . We then compare the values of (14) at these elbow-selected points and retain the pair (K^*, ϕ^*) associated with the minimum. This strategy preserves the role of the information criterion in selecting model complexity, while avoiding the degeneracy of (13), and is consistent with the practical spirit of spatially constrained clustering approaches such as Chavent et al. (2018) and related works (Mattera and Franses, 2023; Cerqueti and Mattera, 2025; Morelli et al., 2025, 2026), where the spatial tuning parameter is calibrated through a trade-off between fit and spatial coherence rather than by a single global minimization criterion. Fig. 3 illustrates how the selection procedure works.

² The magnitude of the spatial regularization parameter ϕ is inherently scale-dependent and influenced by the sample size n , the likelihood specification, the topology of the spatial graph (e.g., density, contiguity, normalization strategies), and the relative scaling between the likelihood and the spatial coherence term. In the proposed penalized likelihood framework, the aggregated log-likelihood contribution typically grows proportionally with n , while the Potts-type spatial penalty depends on the number and structure of neighboring connections induced by the spatial graph. Consequently, the effective impact of ϕ must be interpreted relative to the scale of both components of the objective function rather than through its absolute numerical magnitude. In large spatial datasets, even numerically small values of ϕ may exert substantial regularizing effects by strongly influencing the trade-off between statistical fit and spatial coherence. For this reason, the absolute value of ϕ is not directly comparable across different applications, spatial resolutions, or graph structures.




Fig. 3. Illustrative information-criterion curves for different values of the spatial regularization parameter ϕ . For each curve, the elbow-selected number of clusters is highlighted. The red circle identifies the best solution, that is, the elbow $K(\phi)$ associated with the lowest IC.

3.4. Fast estimation of large SCSAR

In SAR model (4), the log-likelihood for a sample of size n contains the Jacobian adjustment $\log|\mathbf{I}_n - \rho \mathbf{W}|$, which constrains the parameter space for ρ through the non-singularity of $\mathbf{A}(\rho) = \mathbf{I}_n - \rho \mathbf{W}$. An accurate and efficient evaluation of $\log|\mathbf{I}_n - \rho \mathbf{W}|$ is therefore a central computational task in the estimation of the parameters. This applies also to cluster-level estimates involving $\mathbf{A}_k(\rho_k)$. In this case, the classical solution (Ord, 1975) consists of exploiting the spectral identity

$$\log|\mathbf{I}_{n_k} - \rho_k \mathbf{W}_{kk}| = \sum_{i=1}^{n_k} \log(1 - \rho_k \lambda_{ki}), \quad (15)$$

where λ_{ki} are the eigenvalues of the cluster-specific weights matrix $\mathbf{W}_{kk}$. In principle, this approach, also adopted by Cerqueti et al. (2025b), is attractive as it involves a single spectral decomposition per cluster, thus providing an extremely fast evaluation of the log-determinant and its derivatives for any candidate ρ_k , and exact admissible bounds for ρ_k from the extreme eigenvalues.

A well-known drawback of the eigenvalue-based approach is that it becomes computationally inefficient as n grows (Pace, 1997). In our application, which covers the Italian municipalities ($n \approx 8000$), this issue is particularly important because, even after clustering, many cluster sizes n_k remain large. Moreover, the SCSAR algorithm updates the partition iteratively; each reassignment modifies the submatrices $\mathbf{W}_{kk}$ and therefore leads to spectral decompositions at every iteration. Given the problem size, computing full spectral decompositions can be impractical without substantial computational resources. To be more precise, the complexity without clustering is $\mathcal{O}(n^3)$, while under a clusterwise scheme it scales as $\mathcal{O}(n_k^3)$ for each of the K clusters. Moreover, we have also to consider T outer iterations, so that the overall computational burden is even larger. It is clear that as K and n increase, this quickly becomes intractable. Hence, despite its conceptual elegance, the eigenvalue method is not scalable in our setting. This motivates the development of a feasible sparse alternative.

A practical route to improved performance is to exploit sparsity (Pace and Barry, 1997a). Because spatial interaction is local, most entries of the weight matrix are zero, so $\mathbf{W}$ (and its clusterwise blocks) can be treated as sparse. Leveraging sparse linear algebra substantially reduces the computational burden, bringing the per-cluster cost towards $\mathcal{O}(n_k^2)$, and makes large- n applications far more tractable. In particular, for sparse spatial weights, as in the case of Italian municipalities, exact sparse matrix factorizations provide a powerful alternative (e.g., Pace, 1997; Pace and Barry, 1997b).

To clarify the mathematical argument in the row-standardized case, let $\mathbf{W}_{0,kk}$ denote the symmetric non-standardized contiguity matrix on the subgraph induced by cluster k , and let

$$\mathbf{D}_k = \text{diag}(d_{k1}, \dots, d_{kn_k}), \quad d_{ki} = \sum_{j=1}^{n_k} (\mathbf{W}_{0,kk})_{ij},$$

be the diagonal matrix of within-cluster row sums. The row-standardized cluster-specific matrix is then $\mathbf{W}_{kk} = \mathbf{D}_k^{-1} \mathbf{W}_{0,kk}$, which is in general asymmetric. However, one may define the symmetric matrix

$$\mathbf{W}_{kk}^* = \mathbf{D}_k^{-1/2} \mathbf{W}_{0,kk} \mathbf{D}_k^{-1/2}.$$

Since $\mathbf{W}_{kk}^* = \mathbf{D}_k^{1/2} \mathbf{W}_{kk} \mathbf{D}_k^{-1/2}$, the matrices $\mathbf{W}_{kk}$ and $\mathbf{W}_{kk}^*$ are similar and therefore share the same eigenvalues. Consequently, the corresponding SAR matrices

$$\mathbf{A}_k(\rho_k) = \mathbf{I}_{n_k} - \rho_k \mathbf{W}_{kk} \quad \text{and} \quad \mathbf{A}_k^*(\rho_k) = \mathbf{I}_{n_k} - \rho_k \mathbf{W}_{kk}^*$$

are also similar, because $\mathbf{A}_k^*(\rho_k) = \mathbf{D}_k^{1/2} \mathbf{A}_k(\rho_k) \mathbf{D}_k^{-1/2}$. Hence, $\det(\mathbf{A}_k(\rho_k)) = \det(\mathbf{A}_k^*(\rho_k))$, so that the Jacobian term entering the SAR log-likelihood is preserved exactly under this transformation. This provides the mathematical justification for symmetry-based sparse factorization methods in clustered row-standardized settings.

For a symmetric weight matrix $\mathbf{W}_{kk}$, the sparse Cholesky factorization $\mathbf{A}_k(\rho_k) = \mathbf{L}_k(\rho_k) \mathbf{L}_k(\rho_k)'$ yields (Smirnov and Anselin, 2001)

$$\log |\mathbf{A}_k(\rho_k)| = 2 \sum_{i=1}^{n_k} \log l_{k,ii}(\rho_k). \quad (16)$$

Therefore, in the row-standardized case, the determinant can be evaluated through a similarity-transformed symmetric system without any approximation of the Jacobian term, provided that the transformed matrix is symmetric positive definite. In this sense, the use of sparse factorization methods is not merely a computational shortcut, but is also supported by an exact determinant-preserving transformation in the clustered row-standardized setting.

The motivation of using factorization approaches is the following. Since the zero pattern of $\mathbf{A}_k(\rho_k) = \mathbf{I}_{n_k} - \rho_k \mathbf{W}_{kk}$ does not change with ρ_k , within each estimation step, taking the clusters fixed, we perform a single symbolic factorization per cluster while requiring only the inexpensive numeric re-factorizations for each candidate ρ_k explored by the maximization step. Therefore, the computational cost scales with the number of non-zero elements. In the case of municipal data, where each row has few neighbors, this yields a considerable improvement in computational time. Most importantly, when labels change, only those clusters whose index sets $\mathcal{I}_k$ have been modified require a new symbolic phase for detecting the matrix structure across the SCSAR iterations. Unchanged clusters use permutations and data structures already computed. By contrast, eigenvalue-based schemes typically require recomputing spectra whenever the clusters change, which is costly computationally. We implement this new feature in the R package SCDA (Cerqueti et al., 2025b).

The computational gain can be understood more formally by contrasting one outer iteration of a standard eigenvalue-based SCSAR implementation with one outer iteration of the sparse implementation. Let n be the total number of spatial units, K the number of clusters, and n_k the size of cluster k , with $\sum_{k=1}^K n_k = n$. Moreover, let $m = \text{nnz}(\mathbf{W})$ denote the number of non-zero entries of the sparse spatial weights matrix. In an eigenvalue-based implementation, the dominant cost at each outer iteration is the repeated spectral decomposition of the cluster-specific matrices $\mathbf{W}_{kk}$, yielding a complexity of order $O\left(\sum_{k=1}^K n_k^3\right)$. Adding the assignment step, one outer iteration has overall cost $O\left(\sum_{k=1}^K n_k^3 + mK + nK\right)$. By contrast, in the sparse implementation the spectral decomposition is replaced by sparse factorization and sparse determinant evaluation. Following the computational spatial econometrics literature (Pace and Barry, 1997b; Smirnov and Anselin, 2001), the estimation cost is then substantially reduced and may be represented as $O\left(\sum_{k=1}^K n_k^2\right)$, so that the overall cost per outer iteration becomes $O\left(\sum_{k=1}^K n_k^2 + mK + nK\right)$. Notably, this expression shows that the proposed fast implementation does not remove the iterative nature of SCSAR; rather, it reduces the dominant cost of each estimation step by replacing repeated spectral decompositions with sparse factorization-based determinant evaluations.

Under approximately balanced clusters, that is $n_k \approx n/K$, the above cost functions become

$$O\left(\frac{n^3}{K^2} + nK\right)$$

for the eigenvalue-based approach and

$$O\left(\frac{n^2}{K} + nK\right)$$

for the sparse implementation. These expressions also clarify how the computational burden changes with the main structural dimensions of the problem. First, as n increases, both implementations become more expensive, but at different rates. Under approximately balanced clusters, the dominant estimation cost grows as $O(n^3/K^2)$ in the eigenvalue-based approach and as $O(n^2/K)$ in the sparse implementation. Therefore, the computational advantage of sparse factorization becomes more pronounced as the number of spatial units increases. Second, the number of clusters K has an ambiguous effect. On the one hand, larger values of K reduce the size of the cluster-specific blocks, which is favorable for the estimation step. On the other hand, they increase the cost of the assignment step and the overall bookkeeping burden, since each unit must be compared against a larger number of candidate groups. Third, cluster imbalance also matters. Since both $\sum_{k=1}^K n_k^3$ and $\sum_{k=1}^K n_k^2$ are minimized when clusters are balanced, highly unbalanced partitions are computationally less favorable, especially under the eigenvalue-based strategy.

4. Computational benchmarking

In this section, we compare the standard eigenvalue-based implementation and the proposed sparse implementation on spatial subsamples extracted from the Italian municipal geometry. The purpose of this benchmark is not merely to confirm that sparse factorization is faster than eigendecomposition in large spatial autoregressive models, a result that is already well understood in the literature. Rather, the relevant question in the present clustered framework is how the computational burden changes jointly with the number of spatial units n and the number of clusters K . This issue is not trivial: on the one hand, increasing n makes the estimation problem more demanding for any numerical implementation; on the other hand, increasing K splits the global spatial system into smaller cluster-specific subproblems and may therefore reduce the computational burden of clusterwise estimation, at least when the resulting clusters are not excessively unbalanced. Hence, in the clustered SCSAR framework, the overall computational cost is governed by two opposing forces: larger spatial systems tend to slow down estimation, whereas finer partitions may speed it up by




Fig. 4. Visual representation of the computational benchmark design. Starting from the full Italian municipal geometry, a seed municipality is selected. Around that point, increasingly large spatial subsets are constructed by selecting the nearest municipalities in centroid distance. On each subset, the same SCSAR model is repeatedly estimated for different values of K and different numerical implementations.

reducing the size of the matrices involved in each cluster-level fit. The aim of the benchmark is therefore to study how these two dimensions interact in practice, and how the eigenvalue-based and sparse implementations behave as (n, K) vary.

To address this issue, we construct the benchmark directly on the Italian municipal geometry. Rather than relying on artificial Monte Carlo layouts, we progressively enlarge real spatial subsamples extracted from the original map of Italy. The logic of the procedure is illustrated in Fig. 4. We first choose a municipality as a starting point and then build larger and larger spatial subsets by including the geographically nearest municipalities around it. In this way, for each target size n , we obtain a spatially coherent portion of the Italian territory, preserving the actual geographical arrangement and local contiguity structure of the data. This construction is repeated for different starting municipalities. By considering multiple starting points, we verify that the benchmark results do not depend on a specific initialization or on one particular portion of the map, but instead reflect a stable computational pattern across different spatial configurations. The selected ten random initial points c_0 are shown in Fig. 5. The reported timings are therefore averaged over these alternative starting locations, so as to provide a more robust picture of the relative performance of the competing implementations.

For each spatial subset, we estimate the same SCSAR specification under different values of K , keeping the remaining modeling choices fixed. In particular, we set $\phi = 0.1$, that is, a small positive value that avoids excessively rigid and almost purely spatial partitions while leaving the computational comparison essentially unaffected. The benchmark is thus designed to isolate the role of the two main structural dimensions of the problem, namely the number of spatial units and the number of clusters, and to compare how the `eigen` and `Matrix` implementations react to changes in these two quantities. For each benchmark run, we record the total execution time, the total time spent in the outer loop, the average time per outer iteration, the average time spent in the estimation step, the average time spent in the label-updating step, the final refit time, the number of completed iterations, and the convergence status. Among these quantities, the most relevant summary measure for practical purposes is the total computational time, since it directly captures the overall feasibility of the estimator in large spatial systems. In the empirical benchmark reported below, the set of numerical implementations is given by `Base` and `Fast`. Here, `Base` denotes the standard eigenvalue-based implementation, whereas `Fast` denotes the sparse factorization approach. The benchmark is carried out over the sample-size grid $\mathcal{N} = \{1000, 1292, 1670, 2159, 2790, 3605, 4659, 6022, 7782\}$ and the cluster grid $\mathcal{K} = \{2, 3, 4, 5\}$. For each pair (n, K) and each method, the benchmark is repeated across ten distinct seed municipalities, and the reported timing measures are averaged over these alternative starting configurations.

Fig. 6 reports the log mean total computational time for the baseline eigenvalue-based implementation and the fast sparse implementation, distinguishing the different values of K , while Fig. 7 reports the corresponding speed-up ratio, defined as the ratio between the mean total time of the baseline implementation and that of the fast implementation. Two main results emerge. First, the baseline eigenvalue-based implementation displays a clear increase in computational time as n grows, especially for smaller values of K . This behavior is consistent with the high cost of repeated spectral decompositions on large cluster-specific spatial matrices. Increasing K mitigates this burden, since the original spatial system is split into smaller cluster-level subproblems. Second, the fast sparse implementation remains much more stable across the benchmark grid. Its execution time is not strictly monotone in n , which is expected because sparse-factorization runtimes depend not only on the number of spatial units, but also on the induced graph structure and the balance of the estimated clusters. Fig. 7 confirms this interpretation. Indeed, the fast implementation is almost always faster than the baseline implementation, and its advantage becomes substantial for larger spatial systems. The gain is especially pronounced for small K , where the eigenvalue-based method still has to handle relatively large cluster-specific matrices (in the largest configurations, the advantage is substantial, reaching more than one hundred times for $K = 2$). As K increases, the baseline method also benefits from the partitioning of the system, and the speed-up ratio decreases.



Fig. 5. Selected seed municipalities c_0 for the computational experiment.




Fig. 6. Log mean total computational time as a function of the number of spatial units n , by numerical implementation and number of clusters. Results are averaged across ten seed municipalities shown in Fig. 5.

5. Empirical results and discussion

In what follows, we discuss the results obtained using both a single global model and different cluster-specific models. We compare four specifications considering alternative dependent variables, corresponding to electoral concentration measured by the




Fig. 7. Speed-up ratio of the fast sparse implementation relative to the baseline eigenvalue-based implementation, computed as baseline time divided by fast time. Values larger than one indicate an advantage of the fast implementation. Results are averaged across ten seed municipalities shown in Fig. 5.

Gini and polarized Shannon indices, separately for the Chamber of Deputies and the Senate. Finally, to assess the robustness of the findings, we compare the obtained results with region-specific (i.e., group-fixed) models in terms of goodness-of-fit. For estimation, all the variables are standardized in all model specifications, so that the estimated coefficients are directly comparable across covariates, clusters and alternative specifications themselves.

5.1. Global model

The estimates from the global SAR models are reported in Table 3. Results are generally consistent across the Chamber and the Senate, with most covariates displaying the same sign even though magnitudes vary. Below, we discuss the findings according to the three categories of covariates presented in Table 1.

Results from the global SAR models indicate that certain determinants of electoral concentration are robust across specifications, whereas others vary depending on the measure used. In particular, average income (AvgInc) emerges as a robust negative determinant, indicating that wealthier municipalities are associated with more fragmented competition. A similar pattern is observed for the share of foreign residents (FPop), population density (POPd), the population dependency index (PDI), thus suggesting that demographic factors also contribute significantly. By contrast, industrial land use (LUind) and the mayor's level of education (Degree) show positive and statistically significant coefficients in all specifications. This indicates that municipalities characterized by higher industrial land use and by more educated mayors tend to display higher electoral concentration, consistently across the four outcomes. The share of workers in low-productivity sectors (WLP) is generally negative, although the effect is not significant in the Gini-Senate specification. The share of poverty (SoP) is also negative in all models, but its effect is stronger and more stable for the Gini-based specifications than for the Shannon-based ones.

Finally, the spatial autoregressive parameter ρ is positive, highly significant, and large in all specifications (ranging from 0.63 to 0.70). This indicates strong and persistent spatial spillovers: the level of electoral concentration in a municipality does not depend solely on its internal characteristics but is also influenced by the behavior of neighboring municipalities. In substantive terms, electoral outcomes in Italy exhibit a clear territorial contagion dimension, whereby local concentration patterns tend to mirror those of adjacent territories. However, such a large global ρ cannot reveal from which areas these spillovers originate, nor whether they are homogeneous across the country. This limitation motivates the adoption of clustered spatial models, which can uncover distinct territorial regimes and provide a more nuanced understanding of spatial heterogeneity in electoral concentration. In this regard, we also remark that the goodness-of-fit measures in Table 3 indicate that the explanatory power of the global SAR model is limited.

Table 3
Global SAR models with alternative dependent variables and goodness-of-fit measures.

	Gini Chamber	Gini Senate	Shannon Chamber	Shannon Senate
(Intercept)	−0.0196* (0.0077)	−0.0233** (0.0077)	−0.0088 (0.0080)	−0.0127 (0.0072)
AvgInc	−0.1296*** (0.0183)	−0.1391*** (0.0172)	−0.1353*** (0.0176)	−0.1414*** (0.0169)
Gini	−0.0123 (0.0094)	−0.0099 (0.0092)	−0.0154 (0.0082)	−0.0129 (0.0078)
FPop	−0.0556*** (0.0079)	−0.0495*** (0.0083)	−0.0730*** (0.0078)	−0.0669*** (0.0080)
SoP	−0.0977*** (0.0190)	−0.1291*** (0.0202)	−0.0459* (0.0224)	−0.0595** (0.0200)
POPd	−0.2188*** (0.0084)	−0.2295*** (0.0084)	−0.1438*** (0.0084)	−0.1494*** (0.0081)
Degree	0.0344*** (0.0073)	0.0391*** (0.0078)	0.0166* (0.0078)	0.0205** (0.0076)
Age	−0.0018 (0.0030)	−0.0038 (0.0058)	−0.0043 (0.0171)	−0.0058 (0.0073)
ESI	−0.0056 (0.0071)	−0.0170 (0.0089)	0.0056 (0.0138)	−0.0024 (0.0075)
PDI	−0.1376*** (0.0080)	−0.1365*** (0.0082)	−0.1442*** (0.0081)	−0.1447*** (0.0079)
ER	−0.0664*** (0.0153)	−0.0853*** (0.0143)	−0.0100 (0.0173)	−0.0169 (0.0138)
LUind	0.0763*** (0.0087)	0.0876*** (0.0086)	0.0360*** (0.0086)	0.0455*** (0.0083)
WLP	−0.0380*** (0.0097)	−0.0025 (0.0097)	−0.0506*** (0.0096)	−0.0267 (0.0097)
ρ	0.6536*** (0.0093)	0.6315*** (0.0098)	0.6968*** (0.0087)	0.6874*** (0.0090)
R ²	0.1342	0.1317	0.1253	0.1230
RMSE	5.3418	3.5346	7.5402	35.5957
MAE	3.6123	2.2830	4.8476	15.4073
Num. obs.	7782	7782	7782	7782
LR test	3251.5085***	2826.8230***	3905.4752***	3624.8267***

Note: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

5.2. Clustered models

In the case of the SCSAR model, the joint selection of the number of clusters K and the spatial penalty parameter ϕ is carried out using the information criterion in (14). Rather than fixing ϕ a priori, we consider a finite grid of candidate values scaled by the sample size, namely $\phi \in \left\{0, \frac{1}{n}, \frac{10}{n}, \frac{50}{n}, \frac{100}{n}, \frac{500}{n}, \frac{1000}{n}\right\}$. Therefore, we do not consider larger values than 0.128 because excessively large penalties tend to over-emphasize spatial contiguity and may deteriorate the information criterion by forcing overly rigid spatial partitions. We remark that the value of ϕ leading to such an effect is substantially lower in the large n setting compared to low dimensional ones (such as in Cerqueti et al., 2025b; Sugawara and Murakami, 2021). The tuning of (K, ϕ) follows the sequential procedure described in Section 3.2 and Fig. 8 summarizes this procedure for all empirical specifications.

Since we consider two concentration measures, namely the Gini index and the normalized Shannon entropy, and two electoral settings, namely the Chamber of Deputies and the Senate, the figure reports a total of four model specifications. The first and second columns report the results for the Chamber of Deputies and the Senate, respectively, using the Gini index, while the third and fourth columns present the corresponding results based on the Shannon index. For each specification, two panels are shown: one for the selection of K conditional on the selected ϕ^* , and one for the selection of ϕ based on the criterion evaluated at the elbow-selected number of clusters K^* .

The results show that in both the Gini-based and Shannon-based specifications, the Chamber of Deputies exhibits a more fragmented clustering structure than the Senate. In particular, the selected number of clusters is larger for the Chamber, with $K = 7$ and $K = 6$ groups, whereas in the Senate the selected solution involves $K = 3$ or $K = 5$ groups depending on the considered dependent variable. This indicates that the spatial heterogeneity of Italian municipalities is systematically stronger in the Chamber-based models than in the Senate-based ones. A second relevant result is that, in all specifications, the selected spatial penalty ϕ is strictly positive, suggesting that spatial regularization plays an important role in the clustered model, improving the information criterion relative to the unpenalized case and supporting the view that spatial contiguity is an empirically relevant component in the identification of interpretable clusters.

Fig. 9 reports the spatial clustering structures obtained for the four outcome variables. The results reveal a marked territorial fragmentation of Italian municipalities, but the degree and spatial organization of this fragmentation vary across electoral arenas and concentration measure. For the Chamber of Deputies, the Gini-based specification in panel (a) identifies a highly fragmented spatial structure. The Shannon-based specification in panel (c) displays a similarly fragmented structure. The Senate results show





Fig. 8. Sequential tuning of the clustered SCSAR models for the four dependent variables considered in the empirical analysis. First column: Chamber of Deputies with Gini index; second column: Senate with Gini index; third column: Chamber of Deputies with Shannon index; fourth column: Senate with Shannon index. For each specification, the first-row panel reports the selection of the number of clusters from $IC(K | \phi^*)$, conditional on the selected value ϕ^* , while the second-row panel reports the selection of the spatial penalty parameter from $IC(K_{elbow}(\phi) | \phi)$, that is, the information criterion evaluated at the elbow-selected number of clusters for each candidate value of ϕ .


a more spatially structured pattern. In the Gini-based specification in panel (b), the selected partition contains only three clusters, with a clearer separation between a Southern cluster, a Centre-Northern cluster, and a mainly Northern cluster. The Shannon-based Senate specification in panel (d) also yields a more fragmented partition into five groups. Therefore, the Senate maps exhibit more territorial compactness than the Chamber maps, indicating that electoral concentration in the Senate is organized around fewer and more spatially coherent regimes.

We remark that the institutional structure of the 2018 Italian general election helps explain the different clustering patterns observed in Fig. 9. Indeed, the Chamber of Deputies was organized into 28 electoral districts, thereby creating a more articulated territorial structure than the Senate, which was instead organized into 20 regional districts. Moreover, all citizens aged 18 and above were eligible to vote for the Chamber of Deputies, whereas Senate voting was restricted to citizens aged 25 and above, implying a younger and potentially more heterogeneous electorate in the Chamber of Deputies elections. These institutional features likely generated higher effective electoral thresholds across several regions, increasing the mechanical and strategic disadvantages faced by smaller parties. For instance, voters face weaker incentives for strategic coordination and may therefore be more willing to support minor parties or locally competitive lists, even when such parties have limited national competitiveness or uncertain probabilities of surpassing the electoral threshold. This mechanism appears to generate a more fragmented and spatially scattered geography of electoral concentration, characterized by a larger number of localized electoral regimes. By contrast, the Senate results suggest a stronger territorial coordination of voting behavior. Indeed, Senate seats were constitutionally allocated on a regional basis, the chamber itself had fewer seats and smaller effective district magnitudes, and the electoral competition was therefore characterized by stronger coordination incentives among parties and voters. In this setting, strategic voting and coalition alignment become more relevant, reducing the electoral viability of smaller parties and favoring broader territorial aggregation patterns. As a result, municipalities tend to cluster into fewer and more spatially coherent electoral regimes.


This interpretation is further supported by the spillover maps, shown in Fig. 10. The Gini-Senate specification in panel (b) is particularly informative from a substantive political perspective. The areas characterized by stronger spillover effects closely resemble some of the most historically consolidated political territories in Italy. In particular, municipalities located in regions such as Emilia-Romagna and Tuscany display highly concentrated patterns, coherently with the long-standing territorial persistence of left-wing electoral dominance. Similarly, strong spillovers emerge in parts of Veneto and Friuli-Venezia Giulia, areas historically associated with cohesive center-right and regionalist voting patterns. These results suggest that electoral concentration in the Senate is strongly shaped by persistent territorial political identities and historically rooted regional voting cultures. By contrast, the Chamber of Deputies spillover maps appear substantially more fragmented and locally heterogeneous. The Gini-Chamber specification in panel (a) highlights several localized pockets of strong concentration, such as Valle d'Aosta and Alto Adige in the North, rather than broad macro-territorial regimes. This pattern is plausibly related to the presence of strong local and autonomist political movements, historically connected to regional identity, linguistic minorities, and territorial autonomy. Interestingly, this pattern becomes less pronounced in the Senate election, suggesting that strategic voting incentives and the more territorially




(a) Clustering structure using $Gini_c$ as dependent variable.



(b) Clustering structure using $Gini_s$ as dependent variable.



(c) Clustering structure using $Shannon_c$ as dependent variable.




(d) Clustering structure using $Shannon_s$ as dependent variable.

Fig. 9. Clustering structure resulting from the SCSAR algorithm using alternative dependent variables and number of clusters.


constrained Senate competition may partially reduce the electoral relevance of smaller local parties. The clusters should not be interpreted as a simple classification of municipalities with high or low spatial autoregressive coefficients. In SCSAR, the partition is determined jointly by the fit of the cluster-specific regression, the local spatial autoregressive structure, and the spatial penalty encouraging neighboring municipalities to share the same label. As a consequence, each cluster represents a local electoral regime in which the relationship between voting concentration, socio-economic covariates and spatial spillovers is internally more homogeneous than in the national model. This interpretation is important for the empirical analysis below, because several covariates change sign or significance across clusters, while the magnitude of ρ also varies across territories and electoral chambers.

The cluster-specific SAR models with the Gini index as dependent variable are shown in [Tables 4](#) and [5](#).


Several covariates emerge as broadly consistent across clusters and chambers, reinforcing the robustness of some patterns already observed in the global models. In particular, average income, population density, the share of foreign residents, and the dependency ratio maintain their expected associations in every cluster of both chambers. In particular, municipalities with higher income levels, denser populations, or more diverse demographics are systematically characterized by more fragmented electoral competition, while a larger dependent population similarly reduces concentration. These patterns confirm that the main socio-economic determinants identified in the global model retain explanatory power even once local heterogeneity is accounted for. Other variables, however, display substantial variation across clusters, underscoring the limits of a global perspective. Employment rate, for example, is strongly negative and significant in some clusters, but turns positive in others, suggesting that its effect on concentration depends heavily on the local economic and political context. Similarly, industrial land use alternates between positive and negative associations, highlighting the differentiated role of local economic structure in shaping party competition. Access to essential services, which appeared irrelevant in the global model, proves significant in most of the clusters, but with diverging signs, that is, in some areas better provision of services reduces concentration, while in others it fosters fragmentation. These results point to the importance of allowing for spatial heterogeneity, since averaging across contexts obscures meaningful local dynamics. Institutional covariates, by




(a) Spatial spillover coefficient considering Gini index as response variable for the Italian Chamber of Deputies election of 2018.



(b) Spatial spillover coefficient considering Gini index as response variable for the Italian Senate election of 2018.



(c) Spatial spillover coefficient considering Shannon index as response variable for the Italian Chamber of Deputies election of 2018.



(d) Spatial spillover coefficient considering Shannon index as response variable for the Italian Senate election of 2018.

Fig. 10. Spatial spillover coefficient ρ resulting from the SCSAR algorithm using alternative dependent variables and number of clusters.

Table 4

Cluster-specific SAR model (Gini Chamber of Deputies).

	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6	Cluster 7
Intercept	−0.3956*** (0.0266)	−0.2222*** (0.0296)	−0.0544*** (0.0034)	−0.1603*** (0.0241)	−1.7351*** (0.3589)	−1.1769*** (0.0540)	0.4716*** (0.0488)
AvgInc	−0.4221*** (0.0075)	−0.0739*** (0.0075)	−0.3336*** (0.0099)	−0.1676*** (0.0056)	−0.0569 (0.0835)	−0.1051*** (0.0145)	−0.0351** (0.0126)
Gini	0.0267*** (0.0039)	0.0093** (0.0033)	0.0105** (0.0034)	0.0080** (0.0025)	0.4079*** (0.0617)	0.0450*** (0.0083)	−0.0226*** (0.0051)
FPop	−0.1313*** (0.0033)	−0.0961*** (0.0038)	−0.2484*** (0.0041)	−0.1100*** (0.0030)	−0.1307*** (0.0387)	−0.0311*** (0.0069)	−0.0946*** (0.0054)
SoP	−0.5812*** (0.0093)	0.1013*** (0.0093)	−0.1119*** (0.0103)	0.0216** (0.0078)	0.1272 (0.0841)	0.1400*** (0.0164)	0.1427*** (0.0159)
POPd	−0.2981*** (0.0031)	−1.2800*** (0.0089)	−0.5042*** (0.0050)	−0.9107*** (0.0067)	−0.0603** (0.0212)	−1.4906*** (0.0175)	−0.4425*** (0.0072)
Degree	0.0236*** (0.0030)	0.0224*** (0.0033)	0.0522*** (0.0034)	0.0280*** (0.0027)	0.1405** (0.0437)	0.0262*** (0.0065)	0.0107* (0.0053)

(continued on next page)

Table 4 (continued).

	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6	Cluster 7
Age	−28.7464*** (2.3387)	−20.3032*** (2.5988)	−0.0107*** (0.0011)	−28.2617*** (2.1105)	−54.1574 (31.3078)	−44.8498*** (4.6713)	−18.5371*** (4.3071)
ESI	0.0492*** (0.0032)	0.0941*** (0.0043)	−0.0600*** (0.0039)	0.0415*** (0.0034)	0.0860 (0.0465)	0.1271*** (0.0077)	0.0618*** (0.0062)
PDI	−0.1165*** (0.0034)	−0.0734*** (0.0040)	−0.1833*** (0.0037)	−0.0840*** (0.0034)	−0.1434*** (0.0321)	−0.0874*** (0.0066)	−0.1397*** (0.0062)
ER	−0.1720*** (0.0062)	0.0813*** (0.0071)	−0.4248*** (0.0072)	−0.0345*** (0.0060)	0.0060 (0.0617)	0.0372** (0.0127)	0.0075 (0.0108)
LUind	0.1077*** (0.0033)	−0.0431*** (0.0038)	0.0602*** (0.0039)	−0.0100** (0.0031)	0.0672 (0.0423)	0.0172* (0.0070)	−0.0225*** (0.0064)
WLP	0.0720*** (0.0041)	−0.2100*** (0.0045)	0.0417*** (0.0046)	−0.0502*** (0.0039)	−0.0490 (0.0421)	−0.1026*** (0.0077)	−0.1197*** (0.0077)
ρ	0.0720*** (0.0129)	0.1754*** (0.0075)	0.1556*** (0.0064)	0.1006*** (0.0082)	0.3934*** (0.0264)	0.1470*** (0.0128)	0.1817*** (0.0107)
R ²	0.9427	0.9743	0.9873	0.9628	0.1746	0.9055	0.8768
MSE	0.0090	0.0121	0.0105	0.0092	1.8024	0.0441	0.0376
RMSE	0.0950	0.1101	0.1023	0.0962	1.3425	0.2099	0.1939
MAE	0.0769	0.0915	0.0833	0.0783	1.0034	0.1754	0.1553
Num. obs.	1062	1161	934	1414	738	1109	1364
LR test	29.1103***	447.0196***	451.4295***	140.0937***	183.0444***	118.6595***	255.6985***

Note: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

Table 5

Cluster-specific SAR model (Gini Senate).

	Cluster 1	Cluster 2	Cluster 3
Intercept	−0.2355*** (0.0390)	−0.1475*** (0.0312)	−0.4595*** (0.0309)
AvgInc	−0.1371*** (0.0066)	−0.2059*** (0.0082)	−0.0116 (0.0349)
Gini	−0.0140*** (0.0039)	−0.0199*** (0.0042)	0.0980*** (0.0269)
FPop	−0.0751*** (0.0044)	−0.0201*** (0.0037)	−0.1347*** (0.0248)
SoP	0.0003 (0.0010)	−0.0463*** (0.0108)	0.0536 (0.0553)
POPd	−0.5022*** (0.0084)	−0.2941*** (0.0053)	−0.1790*** (0.0163)
Degree	0.0135** (0.0044)	0.0114*** (0.0034)	0.1184*** (0.0256)
Age	−22.8050*** (3.4161)	−20.7870*** (2.7533)	0.0030 (0.0087)
ESI	−0.0002 (0.0010)	0.0360*** (0.0041)	−0.0304 (0.0282)
PDI	−0.1352*** (0.0051)	−0.0716*** (0.0040)	−0.1349*** (0.0208)
ER	−0.1264*** (0.0077)	0.2141*** (0.0085)	−0.1377*** (0.0402)
LUind	0.0290*** (0.0045)	0.0225*** (0.0040)	0.1181*** (0.0259)
WLP	0.0062 (0.0056)	−0.0357*** (0.0048)	0.0239 (0.0271)
ρ	0.4535*** (0.0098)	0.4962*** (0.0100)	0.4523*** (0.0180)
R ²	0.8695	0.7497	0.2414
MSE	0.0672	0.0450	1.5569
RMSE	0.2591	0.2122	1.2478
MAE	0.2031	0.1718	0.9227
Num. obs.	3178	2810	1794
LR test	1449.6235***	1611.7629***	477.1044***

Note: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

contrast, become considerably more relevant once local heterogeneity is accounted for: the mayor's level of education is positive and significant in every cluster, and the mayor's age is negative and significant in almost all of them, whereas both were weak in the global model. The spatial autoregressive parameter ρ deserves particular attention. In the Gini Chamber model, ρ is positive and highly significant in all seven clusters, but its magnitude varies substantially. It ranges from a relatively weak spillover in Cluster

Table 6
Cluster-specific SAR model (Shannon Chamber of Deputies).

	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
Intercept	−0.0918*** (0.0059)	−0.0749** (0.0232)	−0.4319*** (0.0749)	0.3285*** (0.0290)	−0.5718 (0.3199)	−0.4615*** (0.0257)
AvgInc	−0.1886*** (0.0160)	−0.3702*** (0.0067)	−0.0114 (0.0165)	−0.3895*** (0.0071)	−0.0259 (0.0689)	−0.2839*** (0.0075)
Gini	−0.0910*** (0.0063)	−0.0300*** (0.0028)	0.0985*** (0.0100)	−0.0329*** (0.0032)	0.0345 (0.0353)	−0.0382*** (0.0034)
FPop	−0.2308*** (0.0065)	−0.1194*** (0.0028)	−0.0939*** (0.0091)	−0.1066*** (0.0035)	−0.1274*** (0.0371)	−0.1007*** (0.0031)
SoP	0.0578*** (0.0167)	−0.2948*** (0.0084)	0.4449*** (0.0230)	−0.2751*** (0.0095)	0.1922** (0.0738)	−0.2116*** (0.0089)
POPd	−0.3405*** (0.0106)	−0.4663*** (0.0051)	−0.9473*** (0.0183)	−0.4537*** (0.0073)	−0.0896*** (0.0183)	−0.5226*** (0.0070)
Degree	0.0652*** (0.0059)	0.0089*** (0.0026)	−0.0596*** (0.0088)	0.0146*** (0.0032)	0.1000** (0.0367)	0.0300*** (0.0029)
Age	−0.0131*** (0.0023)	−10.1758*** (2.0405)	−20.9636*** (6.5724)	−7.9897** (2.5113)	−41.5587 (27.9587)	−10.8362*** (2.2484)
ESI	−0.0490*** (0.0063)	−0.0457*** (0.0033)	0.3162*** (0.0106)	−0.0253*** (0.0039)	−0.0424 (0.0374)	−0.0239*** (0.0038)
PDI	−0.1186*** (0.0067)	−0.1788*** (0.0031)	−0.1397*** (0.0090)	−0.1714*** (0.0036)	−0.1855*** (0.0307)	−0.1960*** (0.0037)
ER	−0.5278*** (0.0122)	−0.1479*** (0.0060)	0.4072*** (0.0184)	−0.1190*** (0.0074)	0.0205 (0.0512)	−0.1303*** (0.0064)
LUind	0.0195** (0.0067)	−0.0043 (0.0029)	0.0787*** (0.0096)	−0.0268*** (0.0038)	0.0940* (0.0394)	−0.0086** (0.0032)
WLP	−0.0299*** (0.0078)	−0.0937*** (0.0035)	−0.3970*** (0.0110)	−0.0849*** (0.0045)	0.0705 (0.0407)	−0.1275*** (0.0038)
ρ	0.2437*** (0.0100)	0.0370*** (0.0077)	0.4542*** (0.0111)	0.0612*** (0.0089)	0.6130*** (0.0187)	0.1080*** (0.0090)
R ²	0.9638	0.9503	0.9027	0.9371	0.1832	0.9378
MSE	0.0476	0.0096	0.0981	0.0136	2.7700	0.0113
RMSE	0.2183	0.0979	0.3133	0.1165	1.6643	0.1061
MAE	0.1793	0.0831	0.2524	0.0982	1.3798	0.0893
Num. obs.	1293	1547	1006	1436	1070	1430
LR test	470.6821***	19.2215***	898.5648***	46.9281***	577.9335***	136.3984***

Note: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

1 ($\rho = 0.0720$) to a much stronger value in Cluster 5 ($\rho = 0.3934$), with the remaining clusters mostly lying between approximately 0.10 and 0.18. This indicates that vote concentration in the Chamber of Deputies displays local spatial dependence everywhere, but the intensity of this dependence is far from homogeneous across the territory. The Gini Senate model shows a different pattern. In this case, the number of selected clusters is smaller and the estimated spillover coefficients are remarkably stable across clusters: $\rho = 0.4535$, $\rho = 0.4962$, and $\rho = 0.4523$, all highly significant. Hence, unlike the Chamber specification, where the strength of spatial dependence varies across clusters, the Senate specification points to a more homogeneous and stronger spatial propagation of vote concentration. This suggests that concentration in the Senate election follows broader territorial regularities, whereas concentration in the Chamber election exhibits a more locally differentiated spatial structure. Taken together, the two Gini specifications point to a clear contrast between chambers rather than between concentration measures, since spatial dependence is stronger in the Senate than in the Chamber of Deputies. The Shannon-based specifications, discussed next, serve primarily as a robustness check on this reading. The cluster-specific SAR models using Shannon entropy as the dependent variable are shown in Tables 6 and 7.

The Shannon-based specifications largely confirm the picture obtained with the Gini index. The same four covariates that are robust in the global models retain a consistent sign across regimes: population density and the population dependency index are negative in every cluster of both chambers, and average income and the share of foreign residents are negative in all but one. Wealthier and denser municipalities, and those with larger dependent populations, are thus systematically associated with lower vote concentration, and hence with more fragmented electoral competition, exactly as in the Gini specifications. Also, as in the Gini specifications, other covariates change sign across regimes: the share of poverty, the employment rate, industrial land use and income inequality are all positive in some clusters and negative in others, confirming that their association with electoral fragmentation is highly context-dependent. Interestingly, institutional indicators are more relevant in the clustered specifications than in the global model. As for the spatial autoregressive parameter, all the estimated ρ values are positive and statistically significant. Their average across regimes is about 0.25 for the Chamber and 0.30 for the Senate, so that the Senate again displays somewhat stronger spatial propagation, in line with the Gini results.

5.3. Region-specific models

To further assess model fit, and in response to the need for comparing the proposed clustered specification with simpler territorial alternatives, we also estimate region-specific SAR models using the administrative regions depicted in Fig. 11. Fig. 12 reports the

Table 7
Cluster-specific SAR model (Shannon Senate).

	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5
Intercept	0.0805*** (0.0053)	−0.4456*** (0.0283)	−0.1778*** (0.0215)	−0.0642 (0.1275)	−0.2421*** (0.0431)
AvgInc	−0.6085*** (0.0137)	−0.1806*** (0.0078)	−0.2681*** (0.0059)	0.0306 (0.0514)	−0.1907*** (0.0095)
Gini	0.0539*** (0.0056)	−0.0323*** (0.0034)	−0.0407*** (0.0028)	0.0530* (0.0222)	−0.0584*** (0.0057)
FPop	−0.2444*** (0.0058)	−0.0549*** (0.0036)	−0.0194*** (0.0026)	−0.1477*** (0.0245)	0.0676*** (0.0049)
SoP	−0.7117*** (0.0157)	0.1739*** (0.0094)	−0.0659*** (0.0077)	0.2234*** (0.0528)	0.2208*** (0.0139)
POPd	−0.2234*** (0.0053)	−0.8470*** (0.0096)	−0.4170*** (0.0051)	−0.1033*** (0.0155)	−0.7154*** (0.0107)
Degree	−0.2076*** (0.0053)	0.1041*** (0.0033)	0.0519*** (0.0024)	0.1006*** (0.0236)	0.0846*** (0.0047)
Age	−0.0194*** (0.0021)	−28.1078*** (2.4840)	−12.6282*** (1.8880)	4.7419 (11.1597)	−11.0623*** (3.7626)
ESI	−0.1163*** (0.0063)	−0.0034 (0.0039)	0.0625*** (0.0027)	0.0080 (0.0254)	0.0886*** (0.0059)
PDI	−0.1609*** (0.0058)	−0.2679*** (0.0039)	−0.1331*** (0.0029)	−0.1824*** (0.0211)	−0.0771*** (0.0052)
ER	−0.4457*** (0.0110)	−0.2395*** (0.0071)	0.1342*** (0.0059)	−0.0465 (0.0356)	0.5105*** (0.0121)
LUind	0.0522*** (0.0057)	−0.0592*** (0.0037)	−0.0331*** (0.0028)	0.1061*** (0.0250)	−0.0410*** (0.0055)
WLP	0.1128*** (0.0070)	−0.0661*** (0.0044)	−0.0355*** (0.0033)	−0.0009 (0.0042)	−0.1562*** (0.0063)
ρ	0.2441*** (0.0104)	0.1229*** (0.0073)	0.1737*** (0.0098)	0.6169*** (0.0142)	0.3152*** (0.0109)
R^2	0.9358	0.9751	0.9191	0.3036	0.9099
MSE	0.0377	0.0150	0.0088	1.9758	0.0379
RMSE	0.1940	0.1225	0.0938	1.4056	0.1946
MAE	0.1596	0.1016	0.0776	1.1145	0.1567
Num. obs.	1322	1504	1601	1851	1504
LR test	452.1937***	256.5680***	281.5150***	1041.0587***	638.9586***

Note: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

comparison between the clustered model and the region-specific SAR models in terms of R^2 . The clustered specification clearly provides a better fit in all four cases. For the Gini-based outcomes, the clustered model achieves $R^2 = 0.810$ for the Chamber of Deputies and $R^2 = 0.597$ for the Senate, compared with average region-specific values of 0.421 and 0.432, respectively. The same qualitative results are obtained using the Shannon-based outcomes. This evidence suggests that the data-driven clustering structure captures territorial heterogeneity more effectively than administrative regional partitions. In other words, electoral regimes do not appear to coincide mechanically with regional boundaries. Some regions achieve relatively good fit, but the overall pattern indicates that the SCSAR partition provides a more flexible and informative representation of spatial heterogeneity. This is consistent with the main motivation of the proposed approach, that is, spatially clustered models can identify latent territorial regimes that may cut across formal administrative units. The comparison with the global SAR model further confirms this interpretation. The global specification provides a limited fit, with pseudo- R^2 values between 0.12 and 0.14 across the four outcomes. This indicates that a single national regression structure is not sufficient to capture the spatial and socioeconomic heterogeneity of electoral concentration across Italian municipalities. By contrast, the clustered specification substantially improves fit while retaining a parsimonious number of spatial regimes.

5.4. Sensitivity to the selection of ϕ

As discussed in Section 3.3, the calibration of (K, ϕ) follows a heuristic rule rather than a procedure with formal guarantees. A natural question is therefore how much the empirical conclusions depend on the particular value ϕ^* returned by that rule, and whether neighboring values of the spatial penalty would lead to materially different results. To address this, we hold K fixed at the value selected for each specification and re-estimate the model over a grid of perturbations of ϕ^* . Since the magnitude of ϕ is scale-dependent — it depends on n , on the topology of the spatial graph and on the relative scaling of the likelihood and penalty components, as discussed in Section 3.3 — the perturbation is applied multiplicatively, setting $\phi = (1 + \eta)\phi^*$ with $\eta \in \{-0.10, -0.05, -0.01, 0, 0.01, 0.05, 0.10\}$. We then compare the resulting solutions with the selected one along two dimensions. The first is the partition, summarized by the adjusted Rand index (ARI, Hubert and Arabie, 1985) of each solution with respect to the one obtained at $\eta = 0$. The second is the vector of estimated spatial parameters. Here a cluster-by-cluster comparison is not meaningful, since cluster labels are only identified up to a permutation and solutions obtained under different values of ϕ




Fig. 11. Administrative regional partition of Italian municipalities.




Fig. 12. Comparison of goodness-of-fit between the clustered SCSAR model and region-specific SAR models. The gray bar reports the R^2 of the clustered model, whereas the colored bars report the R^2 obtained by estimating separate SAR models within each administrative region. Panels refer to the Gini and Shannon outcomes for the Chamber of Deputies and the Senate.




Fig. 13. Sensitivity of the SCSAR solutions to the spatial penalty. For each specification, K is held fixed at its selected value and ϕ is perturbed multiplicatively around ϕ^* (dashed line). Top row: average of the estimated ρ_k across clusters, with a band of one standard deviation across clusters. Bottom row: adjusted Rand index of each partition with respect to the one selected at $\eta = 0$.

need not order the regimes in the same way. We therefore report two label-invariant summaries, that is, the average of $\hat{\rho}_k$ across the K clusters, which measures the overall strength of spatial dependence, and its standard deviation, which measures the degree of spatial heterogeneity that the model detects. Two solutions delivering the same average and the same dispersion describe the same spatial structure, irrespective of how the regimes happen to be labeled. Fig. 13 shows that the estimated spatial structure is essentially unaffected by the perturbation. In all four specifications, both the average of $\hat{\rho}_k$ and its dispersion across clusters are stable over the whole range considered, varying by at most 0.02 between the extreme perturbations. Therefore, the overall level of spatial dependence and the amount of heterogeneity detected between regimes are essentially the same at $\eta = \pm 0.10$ as at ϕ^* . More precisely, for the Shannon index of the Chamber of Deputies the selected partition is exactly reproduced over almost the entire grid, from -10% to $+5\%$, and the ARI falls to 0.79 only at $+10\%$. For the Gini index of the Senate the partition is also exactly reproduced for perturbations up to $\pm 5\%$, with the index falling to about 0.84 at $\pm 10\%$. In the remaining two specifications small reallocations already occur at $\pm 1\%$, with the ARI remaining between 0.76 and 0.97. These results suggest that the qualitative conclusions remain unchanged considering relatively small perturbations of the spatial penalty ϕ .

6. Conclusions

This paper investigates the determinants of voting concentration in Italian municipalities using the spatially-clustered spatial autoregressive (SCSAR) model, which has been adapted to be applied to a large set of spatial units. The contribution of the paper is threefold. First, from a substantive perspective, we provide new evidence that voting concentration across Italian municipalities is characterized not only by spatial dependence, but also by substantial spatial heterogeneity. In other words, the intensity and structure of voting concentration vary across different areas of the country, revealing heterogeneous local patterns in electoral behavior. Second, from an applied methodological perspective, we show how the recently proposed SCSAR framework can be effectively used to study fine-resolution electoral geography, where both covariate effects and spatial spillover intensity may vary across groups of municipalities. Third, from a computational perspective, we make this framework operational on a large municipal dataset by exploiting sparse matrix methods and by documenting, through a dedicated benchmark, the computational gains relative to the standard eigenvalue-based implementation.

The empirical results, based on the 2018 elections, reveal the existence of significant spatial spillovers in the Italian political system. These territorial differences make the spatial clustering approach particularly suitable, allowing for the identification of groups of municipalities with similar structural relationships and spillover effects. In particular, the main factors determining voting concentration are the economic and demographic variables, which prove to be the most robust predictors across all model specifications. Among the others, electoral concentration is influenced by structural factors such as income, demographic composition, and population density. Regarding the spatial spillover, we find heterogeneous but positive spillovers and a larger spatial fragmentation in Chamber of Deputies compared to Senate, coherently with the different electorate composition, district structure, and seat allocation mechanisms of the two chambers. In light of these spillover effects, electoral behavior cannot be

understood solely through the analysis of the characteristics of individual municipalities. As a robustness check, we also compared the spatially-clustered outcome with those obtained by fixed-groups spatial autoregression models in which the Italian administrative regions are used as clusters, as well as with a single global model. Although some individual regions exhibit relatively good fit, the proposed data-driven clustering structure captures territorial heterogeneity much more effectively than both the global SAR model and restricted region-specific specifications based on predefined administrative partitions. Moreover, given the results at the cluster level, the aggregate estimates for the entire country appeared unreasonably large, suggesting that their intensities varies depending on the cluster to which they belong, and are not uniformly high across all of Italy. While the results are similar across both the Gini and Shannon entropy measures of concentration, there are some differences. For example, the number of clusters identified varied depending on the measurement used.

This study has some methodological limitations that should be acknowledged. First, the SCSAR algorithm is iterative and depends on the initialization of the cluster labels. In the empirical analysis, we mitigate this issue by adopting a mixed initialization strategy that combines spatial information and covariate similarity. Nevertheless, future work should investigate more robust initialization schemes and assess the stability of the output. Second, the theoretical guarantees available for the proposed procedure are conditional in nature. For a fixed oracle partition, the cluster-specific estimators inherit the standard SAR asymptotics (Lee, 2004). Beyond this, formal results on the recovery of the spatial regimes, on the selection of their number and on inference accounting for partition uncertainty are not currently available for this class of methods in a single cross-section. Developing an asymptotic theory for the joint estimation of regimes and parameters, together with post-clustering inference for the cluster-specific coefficients, is therefore an important direction for future research. Third, and relatedly, the SAR likelihood is joint and does not exactly factorize across units. Following Cerqueti et al. (2025b), we split the log-determinant into equal per-unit shares, which yields a computationally convenient surrogate of the penalized objective and is what makes the reallocation step feasible at the scale of our application. This is one possible construction among others: alternative surrogates, or alternative computational treatments of the Jacobian term, may well have better properties. What is currently unknown, for any such construction, is how the surrogate relates to the exact objective, and how this relationship affects the selected partition and the calibration of the hyperparameters. Comparing alternative surrogates and characterizing their relationship with the exact criterion is, in our view, the most valuable methodological extension of this work. Fourth, we remark that the present SCSAR framework uses hard cluster assignments, that is, each spatial unit is assigned to one and only one cluster. While this improves interpretability, it ignores uncertainty in regime membership. Future research could develop fuzzy or probabilistic SCSAR models, following the spirit of fuzzy clustered coefficient estimation in the SCR literature (Sugasawa and Murakami, 2021), so that units located near regime boundaries or in transition areas may belong to multiple clusters with different membership degrees. Finally, future work could compare the soft spatial regularization adopted in SCSAR with hard-contiguity approaches such as SKATER-regression (Anselin and Amaral, 2024). In the present paper, spatial coherence is controlled adaptively through the penalty parameter ϕ , whereas SKATER-like approaches impose connected regimes by construction. However, such approaches also raise specific computational challenges in large spatial systems, since the regime search is based on minimum-spanning-tree pruning and repeated evaluation of candidate cuts. Therefore, applying SKATER-based spatial regime methods to high-dimensional or very large- n settings may require dedicated acceleration strategies or ad hoc approximations. A systematic comparison between soft-regularized SCSAR and hard-constrained spatial regime models would therefore be a valuable extension.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.spasta.2026.101043>.

References

- Anselin, L., Amaral, P., 2024. Endogenous spatial regimes. *J. Geogr. Syst.* 26 (2), 209–234.
- Arbia, G., Petrarca, F., 2011. Effects of MAUP on spatial econometric models. *Lett. Spat. Resour. Sci.* 4 (3), 173–185.
- Ballatore, R.M., Mariani, V., 2019. Human capital differentials across urban and rural areas in Italy. The role of migrations. *Ital. Econ. J.* 5, 307–324.
- Barone, G., D'Ignazio, A., De Blasio, G., Naticchioni, P., 2016. Mr. Rossi, Mr. Hu and politics. The role of immigration in shaping natives' voting behavior. *J. Public Econ.* 136, 1–13.
- Bazzoli, N., Lello, E., 2022. The neo-populist surge in Italy between territorial and traditional cleavages. *Rural. Sociol.* 87, 662–691.
- Best, R.E., 2013. How party system fragmentation has altered political opposition in established democracies. *Gov. Oppos.* 48 (3), 314–342.
- Bivand, R., Piras, G., 2015. Comparing implementations of estimation methods for spatial econometrics. *J. Stat. Softw.* 63, 1–36.
- Borghesi, C., Chiche, J., Nadal, J.-P., 2012. Between order and disorder: A 'weak law' on recent electoral behavior among urban voters? *PLoS One* 7 (7), e39916.
- Camera dei Deputati e Senato della Repubblica Italiana, 2022. Il sistema di elezione del parlamento nazionale: L'evoluzione normativa e la disciplina vigente. p. 158, URL <https://www.senato.it/service/PDF/PDFServer/BGT/01176918.pdf>.
- Cerqua, A., Di Matteo, D., 2025. The bright side of breaking up: How territorial fragmentation increases political engagement in local communities. *J. Reg. Sci.* 65, 951–964.
- Cerqueti, R., Ausloos, M., 2015. Socio-economical analysis of Italy: The case of hagiotoponym cities. *Soc. Sci. J.* 52 (4), 561–564.
- Cerqueti, R., D'Urso, P., Mattera, R., 2025a. Fuzzy group fixed-effects estimation with spatial clustering. *ASTA Adv. Stat. Anal.* 1–32.
- Cerqueti, R., Maranzano, P., Mattera, R., 2025b. Spatially-clustered spatial autoregressive models with application to agricultural market concentration in Europe. *J. Agric. Biological Environ. Stat.* 1–35.
- Cerqueti, R., Mattera, R., 2025. Measuring unit relevance and stability in hierarchical spatio-temporal clustering. *Spat. Stat.* (ISSN: 2211-6753) 66, 100880. <https://dx.doi.org/10.1016/j.spasta.2025.100880>, URL <https://www.sciencedirect.com/science/article/pii/S2211675325000028>.
- Cerqueti, R., Rotundo, G., Ausloos, M., 2018. Investigating the configurations in cross-shareholding: A joint copula-entropy approach. *Entropy* 20 (2), 134.
- Chavent, M., Kuentz-Simonet, V., Labenne, A., Saracco, J., 2018. ClustGeo: An R package for hierarchical clustering with spatial constraints. *Comput. Statist.* 33 (4), 1799–1822.

- Cheibub, J.A., Moreira, T., Sin, G., Tanabe, K., 2022. Dynamic party system fragmentation. *Elect. Stud.* 76, 102440.
- Coffé, H., Von Schoultz, Å., 2021. How candidate characteristics matter: Candidate profiles, political sophistication, and vote choice. *Politics* 41 (2), 137–155.
- Coleman, S., 2004. The effect of social conformity on collective voting behavior. *Political Anal.* 12 (1), 76–96.
- Deltas, G., 2003. The small-sample bias of the gini coefficient: Results and implications for empirical research. *Rev. Econ. Stat.* 85 (1), 226.
- Dragan, I.-M., Isaic-Maniu, A., 2019. An innovative model of reliability—The pseudo-entropic model. *Entropy* 21 (9), 846.
- Elhorst, J.P., et al., 2014. *Spatial Econometrics: From Cross-Sectional Data to Spatial Panels*. Vol. 479, Springer.
- Fabretti, A., Ausloos, M., 2005. Recurrence plot and recurrence quantification analysis techniques for detecting a critical regime. Examples from financial market indices. *Internat. J. Modern Phys. C* 16 (05), 671–706.
- Faggian, A., Modica, M., Modrego, F., Urso, G., 2021. One country, two populist parties: Voting patterns of the 2018 Italian elections and their determinants. *Reg. Sci. Policy Pr.* 13 (2), 397–414.
- Fernandez-Navia, T., Polo-Muro, E., Tercero-Lucas, D., 2021. Too afraid to vote? The effects of COVID-19 on voting behaviour. *Eur. J. Political Econ.* 69, 102012.
- Gill, J., 2005. An entropy measure of uncertainty in vote choice. *Elect. Stud.* 24 (3), 371–392.
- Gini, C., 1921. Measurement of inequality of incomes. *Econ. J.* 31 (121), 124–126.
- Giommoni, T., Loumeau, G., 2022. Lockdown and voting behaviour: A natural experiment on postponed elections during the COVID-19 pandemic. *Econ. Policy* 37 (111), 547–599.
- Golosov, G.V., 2010. The effective number of parties: A new approach. *Party Politics* 16 (2), 171–192.
- Golosov, G.V., 2017. Party systems, electoral systems, and legislative fragmentation: A cross-national empirical study. *Party Politics* 23 (5), 487–497.
- Goodchild, M.F., 2004. The validity and usefulness of laws in geographic information science and geography. *Ann. Assoc. Am. Geogr.* (ISSN: 0004-5608) 94 (2), 300–303. <http://dx.doi.org/10.1111/j.1467-8306.2004.09402008.x>.
- Gunther, R., Lobo, M.C., Bellucci, P., Lisi, M., 2015. The changing determinants of the vote. In: *Voting in Old and New Democracies*. Routledge, pp. 150–192.
- Guriev, S., Papaioannou, E., 2022. The political economy of populism. *J. Econ. Lit.* 60 (3), 753–832.
- Hooghe, M., Teepe, W., 2007. Party profiles on the web: An analysis of the logfiles of non-partisan interactive political internet sites in the 2003 and 2004 election campaigns in Belgium. *New Media Soc.* 9 (6), 965–985.
- Hubert, L., Arabie, P., 1985. Comparing partitions. *J. Classification* 2 (1), 193–218.
- Jensen, C.D., Lacombe, D.J., McIntyre, S.G., 2013. A Bayesian spatial econometric analysis of the 2010 UK General Election. *Pap. Reg. Sci.* 92 (3), 651–667.
- Krashinsky, H., 2024. Does a voter's decision to sit out an election depend upon where others stand?: Some quasi-experimental evidence on the effect of local concentrations of voters and electoral competitiveness on voting participation. *Elect. Stud.* 92, 102864.
- Laakso, M., Taagepera, R., 1979. “Effective” number of parties: A measure with application to West Europe. *Comp. Political Stud.* 12 (1), 3–27.
- Lacombe, D.J., Shaughnessy, T.M., 2007. Accounting for spatial error correlation in the 2004 presidential popular vote. *Public Financ. Rev.* 35 (4), 480–499.
- Lall, R., Vilalta, D., 2025. Varieties of anxieties: Disaggregating emotion and voting behaviour in the COVID-19 era. *Br. J. Political Sci.* 55, e87.
- Lambiotte, R., Ausloos, M., 2006. On the genre-fixation of music: A percolation approach. *Eur. Phys. J. B-Condensed Matter Complex Syst.* 50 (1–2), 183–188.
- Lee, L.-F., 2004. Asymptotic distributions of quasi-maximum likelihood estimators for spatial autoregressive models. *Econometrica* 72 (6), 1899–1925.
- Manek, A.S., Shenoy, P.D., Mohan, M.C., Venugopal, K., 2017. Aspect term extraction for sentiment analysis in large movie reviews using gini index feature selection method and SVM classifier. *World Wide Web* 20 (2), 135–154.
- Manley, D., 2021. Scale, aggregation, and the modifiable areal unit problem. In: *Handbook of Regional Science*. Springer, pp. 1711–1725.
- Mansley, E., Demšar, U., 2015. Space matters: Geographic variability of electoral turnout determinants in the 2012 London mayoral election. *Elect. Stud.* 40, 322–334.
- Maranzano, P., Monturano, G., Tridico, P., 2025. The geography of the Italian citizenship income: The role of poverty and inequality in determining spatial heterogeneity across the Italian municipalities. *Ital. Econ. J.* 1–48.
- Marmani, S., Ficcadenti, V., Kaur, P., Dhesi, G., 2020. Entropic analysis of votes expressed in Italian elections between 1948 and 2018. *Entropy* 22 (5), 523.
- Maškarinec, P., 2024. Geography of voter turnout in slovak local elections (1994–2018): The effects of size and contagion on local electoral participation. *Trans. GIS* 28 (7), 2113–2133.
- Mattera, R., Franses, P.H., 2023. Are african business cycles synchronized? Evidence from spatio-temporal modeling. *Econ. Model.* (ISSN: 0264-9993) 128, 106485. <http://dx.doi.org/10.1016/j.econmod.2023.106485>, URL <https://www.sciencedirect.com/science/article/pii/S0264999323002973>.
- Meka, E., 2022. How much opposition? Political fragmentation and changes in democracy. *Contemp. Politics* 28 (5), 517–538.
- Monturano, G., Resce, G., Ventura, M., 2025. Short-term impact of financial support to inner areas. *Ital. Econ. J.* 1–40.
- Morelli, C., Boccaletti, S., Maranzano, P., Otto, P., 2025. Multidimensional spatiotemporal clustering – An application to environmental sustainability scores in Europe. *Environmetrics* (ISSN: 1180-4009) 36 (2), e2893. <http://dx.doi.org/10.1002/env.2893>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/env.2893>.
- Morelli, C., Maranzano, P., Otto, P., 2026. Spatiotemporal clustering of GHGs emissions in Europe: Exploring the role of spatial component. *Spat. Stat.* (ISSN: 2211-6753) 100960. <http://dx.doi.org/10.1016/j.spasta.2026.100960>, URL <https://www.sciencedirect.com/science/article/pii/S2211675326000084>.
- Musolino, D., 2018. The north-south divide in Italy: Reality or perception? *Eur. Spat. Res. Policy* 25 (1), 29–53.
- Ord, K., 1975. Estimation methods for models of spatial interaction. *J. Amer. Statist. Assoc.* 70 (349), 120–126.
- Pace, R.K., 1997. Performing large spatial regressions and autoregressions. *Econom. Lett.* 54 (3), 283–291.
- Pace, R.K., Barry, R., 1997a. Quick computation of spatial autoregressive estimators. *Geogr. Anal.* 29 (3), 232–247.
- Pace, R.K., Barry, R., 1997b. Sparse spatial autoregressions. *Statist. Probab. Lett.* 33 (3), 291–297.
- Passarelli, G., 2017. Determinants of preferential voting in Italy: General lessons from a crucial case. *Representation* 53 (2), 167–183.
- Preda, V., Dedu, S., Gheorghe, C., 2015. New classes of Lorenz curves by maximizing Tsallis entropy under mean and Gini equality and inequality constraints. *Phys. A* 436, 925–932.
- Raffinetti, E., Siletti, E., Vernizzi, A., 2015. On the gini coefficient normalization when attributes with negative values are considered. *Stat. Methods Appl.* 24 (3), 507–521.
- Reményi, P., Gekić, H., Bidžan-Gekić, A., Sümeghy, D., 2023. Contextual factors influencing election results in Bosnia and Herzegovina. *Qual. Quant.* 57 (6), 5693–5712.
- Saariama, T., Tukiainen, J., 2016. Local representation and strategic voting: Evidence from electoral boundary reforms. *Eur. J. Political Econ.* 41, 31–45.
- Shin, M.E., Agnew, J., 2007. The geographical dynamics of Italian electoral change, 1987–2001. *Elect. Stud.* 26 (2), 287–302.
- Shin, M., Agnew, J., 2016. Spatial regression for electoral studies: The case of the Italian lega nord. In: *Revitalizing Electoral Geography*. Routledge, pp. 59–74.
- Singer, M., Gershman, Z., 2018. Do changes in district magnitude affect electoral fragmentation? Evidence over time at the district level. *Elect. Stud.* 54, 172–181.
- Smirnov, O., Anselin, L., 2001. Fast maximum likelihood estimation of very large spatial autoregressive models: A characteristic polynomial approach. *Comput. Statist. Data Anal.* 35 (3), 301–319.
- Sugasawa, S., Murakami, D., 2021. Spatially clustered regression. *Spat. Stat.* 44, 100525.
- Taagepera, R., 1999. Supplementing the effective number of parties. *Elect. Stud.* 18 (4), 497–504.
- Tanak, A.K., Borzadaran, G.M., Ahmadi, J., 2017. Maximum tsallis entropy with generalized gini and gini mean difference indices constraints. *Phys. A* 471, 554–560.
- Truglia, F.G., Zeli, A., 2020. Spatial analysis of economic and social determinants of vote: the case of the European parliament and constitutional referendum votes in Italy. *Ital. Political Sci. Review/Rivista Ital. di Sci. Politica* 50 (2), 173–190.

- Valentin, V., Dinas, E., 2024. Does party-system fragmentation affect the quality of democracy? *Br. J. Political Sci.* 54 (1), 152–178.
- Vitale, L., 2025. The geography of public goods: Vote concentration and distributive politics in Italy. Available At SSRN 5333262.
- Walks, A., Twigg-Molecey, A., et al., 2014. Income Inequality and Polarization in Canada's Cities: An Examination and New Form of Measurement. Cities Centre, University of Toronto.
- Zagórski, P., 2022. Too much to choose from? The long-term effects of political fragmentation on electoral turnout. *Politics* 42 (3), 358–375.
- Zhu, A.X., Turner, M., 2022. How is the third law of geography different? *Ann. GIS* (ISSN: 1947-5683) 28 (1), 57–67. <http://dx.doi.org/10.1080/19475683.2022.2026467>.