

Flexible shrinkage in high-dimensional Bayesian spatial autoregressive models

Pfarrhofer, Michael; Piribauer, Philipp

Published in:
Spatial Statistics

DOI:
[10.1016/j.spasta.2018.10.004](https://doi.org/10.1016/j.spasta.2018.10.004)

Published: 01/01/2019

Document Version
Early version, also known as preprint

[Link to publication](#)

Citation for published version (APA):
Pfarrhofer, M., & Piribauer, P. (2019). Flexible shrinkage in high-dimensional Bayesian spatial autoregressive models. *Spatial Statistics*, 29, 109 - 128. <https://doi.org/10.1016/j.spasta.2018.10.004>

Flexible shrinkage in high-dimensional Bayesian spatial autoregressive models

Michael Pfarrhofer*¹ and Philipp Piribauer²

¹WU Vienna University of Economics and Business

²Austrian Institute of Economic Research (WIFO)

12th September 2018

Abstract

This article introduces two absolutely continuous global-local shrinkage priors to enable stochastic variable selection in the context of high-dimensional matrix exponential spatial specifications. Existing approaches as a means to dealing with overparameterization problems in spatial autoregressive specifications typically rely on computationally demanding Bayesian model-averaging techniques. The proposed shrinkage priors can be implemented using Markov chain Monte Carlo methods in a flexible and efficient way. A simulation study is conducted to evaluate the performance of each of the shrinkage priors. Results suggest that they perform particularly well in high-dimensional environments, especially when the number of parameters to estimate exceeds the number of observations. For an empirical illustration we use pan-European regional economic growth data.

Keywords: Matrix exponential spatial specification, model selection, shrinkage priors, hierarchical modeling

JEL Codes: C11, C21, C52

*Corresponding author: Michael Pfarrhofer, Department of Economics, WU Vienna University of Economics and Business. Address: Welthandelsplatz 1, 1020 Vienna, Austria. Email: michael.pfarrhofer@wu.ac.at. The authors would like to thank Maximilian Böck, Niko Hauzenberger, Florian Huber and Gregor Zens for valuable comments and suggestions.

1 Introduction

In the regional economic literature, spatial econometric model specifications have gained momentum in empirical research as a means to explicitly account for spillover effects between geographically structured units. Increasing availability of data often results in (spatial autoregressive) models where the number of observations is relatively small compared to the number of potential covariates. Standard estimation techniques in such environments therefore typically result in imprecise parameter estimates. In the case of severe overparameterization, where the dimensionality of the regressors exceeds the number of observations, standard estimation approaches may even be infeasible.

Bayesian model averaging techniques to alleviate the problems of overparameterization in spatial autoregressive models have been proposed (LeSage and Parent, 2007) and widely applied, particularly in the empirical study of regional economic growth (see, for example, LeSage and Fischer 2008, Crespo Cuaresma and Feldkircher 2013, Crespo Cuaresma et al. 2014, Piribauer and Fischer 2015, or Crespo Cuaresma et al. 2018). These approaches use weighted averages of parameter estimates based on a multitude of potential combinations of explanatory variables, rather than relying on inference based on a single model specification (for extensive discussions see Steel 2017 or Koop 2003). Bayesian model averaging, however, involves the computation of marginal likelihoods for integrating out the underlying model uncertainty. In contrast to classical linear model frameworks, no closed-form solutions for marginal likelihoods in spatial autoregressive specifications are available (LeSage and Parent, 2007), resulting in a severe computational burden especially in high-dimensional estimation problems.

Recent advances in the spatial econometric literature focus on extensions and generalizations of standard spatial autoregressive specifications. For example, some extensions aim at explicitly controlling for unobserved heterogeneity by finite mixture or threshold specifications (Piribauer 2016; Cornwall and Parent 2017), allowing for heterogeneous parameters across space (LeSage and Chih 2017), heteroskedastic specifications of the innovations (LeSage 1997), considering continuous spatial effects (Laurini 2017), or accounting for uncertainty in the spatial weight matrix specification and the underlying nature of spillover processes (LeSage and Pace 2007; Halleck Vega and Elhorst 2015; Piribauer and Fischer 2015), among several others. However, such extensions to the spatial autoregressive modeling framework further increase the emanating computational burden of Bayesian model averaging, rendering the approach computationally intractable.

For spatial autoregressive model specifications, work by Piribauer (2016) and Piribauer and Crespo Cuaresma (2016), for example, uses Bayesian stochastic search variable selection priors (SSVS, George and McCulloch, 1993) as an alternative to Bayesian model averaging. The comparative flexibility of this approach allows to easily extend and generalize the basic framework to more complex specifications.¹ Specifically, the proposed hierarchical modeling approach assumes coefficients of a saturated model under scrutiny to come from a mixture of two Gaussians centered on zero with a spike and a slab component (low and high variance). A binary latent indicator thereby identifies promising subsets of predictors by shrinking unimportant components towards zero. Despite its elegance and appealing theoretical properties, difficulties arise for stochastic search variable selection in large data sets due to the necessity of stochastic search over an enormous space. This implies slow mixing and convergence during estimation (Bhattacharya et al., 2015). Flexible alternatives for shrinkage in high-dimensional econometric frameworks have therefore been advocated (Polson and Scott, 2010; Griffin and Brown, 2017).

¹The computational flexibility of these shrinkage priors comes from the fact that they can easily be implemented in standard Bayesian Markov-chain Monte Carlo (MCMC) algorithms without the need for calculating marginal likelihoods (George and McCulloch, 1993).

In this paper we aim at generalizing variants of the Normal-Gamma (NG, [Griffin and Brown, 2010](#)) and the Dirichlet-Laplace (DL, [Bhattacharya et al., 2015](#)) shrinkage priors to spatial autoregressive specifications. For the purpose of illustrating prior specific properties in the presence of spatial dependence, we carry out an extensive simulation study using synthetic data sets, considering different scenarios of the number of available covariates and degrees of sparsity of the coefficient vector. The paper moreover considers matrix exponential spatial specifications (MESS), introduced by [LeSage and Pace \(2007\)](#) to model global spillover processes. Our results indicate that conventional stochastic search variable selection priors in the spirit of [George and McCulloch \(1993\)](#) work well in relatively low-dimensional settings, where the number of potential covariates do not exceed those of the observations. In high-dimensional modeling frameworks, however, both the Normal-Gamma and Dirichlet-Laplace shrinkage prior exhibit stellar empirical properties. They provide a high degree of adaptiveness of shrinkage, with the Normal-Gamma excelling in terms of precision, while the Dirichlet-Laplace performs particularly well in terms of point estimates.

The remainder of this paper is organized as follows. [Section 2](#) presents the Bayesian spatial econometric framework, followed by the introduction of Normal-Gamma and Dirichlet-Laplace shrinkage priors to spatial autoregressive specifications in [Section 3](#). A simulation study as a means to comparing the properties of the proposed shrinkage priors is presented in [Section 4](#). [Section 5](#) illustrates the performance of the proposed shrinkage priors using pan-European regional economic growth data. [Section 6](#) concludes.

2 Econometric framework

We start by considering a model of the form

$$\mathbf{S}(\bullet)\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}), \quad (1)$$

where \mathbf{y} is an N -dimensional vector of dependent variables and $\mathbf{S}(\bullet)$ describes a linear transformation dependent on a not yet specified parameter. \mathbf{X} is an $N \times K$ matrix of explanatory variables (with a vector of ones in the first column), and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_K)'$ is a K -dimensional vector of parameters. We assume the error term $\boldsymbol{\epsilon}$ to be normally distributed with zero mean and $N \times N$ variance-covariance matrix $\boldsymbol{\Omega}$.

Standard econometric models capturing spatial spillover effects are thoroughly discussed in [LeSage and Pace \(2009\)](#). Conventional spatial autoregressive models (SAR) typically set $\mathbf{S}(\xi) = (\mathbf{I}_N - \xi\mathbf{W})$, where \mathbf{W} is an $N \times N$ exogenous right stochastic spatial weights matrix. If observations $i = 1, \dots, N$ and $j \neq i$ are considered neighbors, then $W_{ij} \neq 0$, otherwise $W_{ij} = 0$. By convention $W_{ii} = 0$, meaning that a region or other spatial unit is not considered a neighbor to itself ([LeSage and Pace, 2009](#)). ξ is a (scalar) spatial parameter with stability condition $|\xi| < 1$. The inverse of $\mathbf{S}(\xi)$ under the specifying assumptions for \mathbf{W} and ξ can be expressed as $(\mathbf{I}_N - \xi\mathbf{W})^{-1} = \sum_{l=0}^{\infty} \xi^l \mathbf{W}^l$, implying global geometric decay over space. An alternative to such a geometric decay pattern is given by the matrix exponential spatial specification (MESS), as proposed by [LeSage and Pace \(2007\)](#), where we set

$$\mathbf{S}(\rho) = \exp(\rho\mathbf{W}) = \sum_{l=0}^{\infty} \rho^l \mathbf{W}^l (l!)^{-1}, \quad (2)$$

with the (scalar) parameter ρ taking the role of measuring spatial dependence. Consequently, these modeling approaches nest the classical linear regression model in the case where the respective spatial

dependence parameter is equal to zero. The main difference between MESS and SAR models is that spatial externalities are modeled by an exponential rather than a geometric decay.²

MESS is advantageous to the conventional SAR approach especially in a Bayesian framework. This is due to the properties of matrix exponentials presented in [LeSage and Pace \(2007\)](#),

- (i) $S(\rho)$ is non-singular,
- (ii) $S(\rho)^{-1} = \exp(\rho W)^{-1} = \exp(-\rho W)$, and
- (iii) $\det(\exp(\rho W)) = \exp(\text{tr}(\rho W))$,

which facilitate the likelihood function given in [Eq. \(3\)](#). Albeit both specifications typically produce rather similar estimates and inference (see, for example, [LeSage and Pace 2007](#), [Piribauer and Fischer 2015](#), or [Strauss et al. 2017](#)), matrix exponential spatial specifications have some potential computational advantages, specifically in high-dimensional data environments. A correspondence between spatial coefficients in the conventional SAR and MESS is stated by [LeSage and Pace \(2007\)](#), where $\xi \approx 1 - \exp(\rho)$. However, it is worth noting that [Debarsy et al. \(2015\)](#) stress that no precise one-to-one correspondence can be established and the two spatial specifications should not be considered perfect substitutes, due to $\rho \in (-\infty, \infty)$ while stability conditions in the conventional SAR case require the parameter space of ξ to be constrained ([LeSage and Pace, 2009](#)). Assuming a normally distributed error term as in [Eq. \(1\)](#), the likelihood of the model is given by

$$p(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}, \sigma^2, \rho) = (2\pi)^{-N/2} \det(\boldsymbol{\Omega})^{-1/2} \exp(-\boldsymbol{\varepsilon}'\boldsymbol{\Omega}^{-1}\boldsymbol{\varepsilon}/2), \quad (3)$$

where we define $\boldsymbol{\varepsilon} = (S(\rho)\mathbf{y} - \mathbf{X}\boldsymbol{\beta})$ for notational convenience.

3 Flexible shrinkage in high-dimensional Bayesian spatial autoregressive models

One conventional approach for cases where the vector of coefficients is expected to be sparse, but without prior knowledge which variables to exclude, is given by penalized least squares. A prominent example is the least absolute shrinkage and selection operator (LASSO) introduced by [Tibshirani \(1996\)](#). This paper stresses the correspondence between the conventional LASSO and a Bayesian approach involving independent double-exponential priors on regression coefficients, a notion further elaborated on in [Park and Casella \(2008\)](#). Given the choice of a specific mixture distribution D , following [Griffin and Brown \(2010\)](#), these priors can typically be expressed as

$$\beta_r | \Psi_r \sim N(0, \Psi_r), \quad \Psi_r \sim D. \quad (4)$$

for $r = 1, \dots, K$. In this setting, the marginal distribution for β_r has heavier than normal tails but places substantial mass on zero. [Griffin and Brown \(2010\)](#) indicate that the standard discrete spike-and-slab prior (where a variant is presented in [Appendix A](#)) may be represented in this form, dependent on the mixture distribution being chosen accordingly.

In addition, various other cases such as the double-exponential mentioned above (leading to the Bayesian LASSO, which is closely tied to the Normal-Gamma prior presented later on), arise for different mixing distributions ([Park and Casella, 2008](#)). Note that in this basic framework, the specific shape of D depends on the deterministic choice of prior hyperparameters. More flexibility and adaptive shrinkage

²It is moreover worth noting that a simple model extension frequently used in empirical applications also incorporates an explicit spatial structure in the exogenous variables by including a spatial lag of \mathbf{X} in [Eq. \(1\)](#). Such a specification is commonly referred to as a spatial Durbin model specification ([LeSage and Pace, 2009](#)).

can be achieved by introducing additional hierarchical layers of priors at this stage. Approaches in this spirit are termed global-local shrinkage priors (Polson and Scott, 2010), and seem to perform well in high-dimensional settings (for instance in time-series analysis, see Bitto and Frühwirth-Schnatter, 2016; Kastner, 2016; Huber and Feldkircher, 2017; Feldkircher et al., 2017). In the following, we discuss two alternatives of continuous global-local shrinkage, the Normal-Gamma and the Dirichlet-Laplace shrinkage prior.

3.1 Normal–Gamma shrinkage prior

A variant of the Normal-Gamma global-local shrinkage prior, as proposed by Griffin and Brown (2010), is given by a scale mixture of Gaussians,

$$\begin{aligned}\beta_r | \psi_r &\sim N(0, 2\lambda^{-2}\psi_r), \\ \psi_r &\sim G(\theta, \theta), \\ \lambda^2 &\sim G(d_0, d_1),\end{aligned}\tag{5}$$

where ψ_r is an idiosyncratic scaling parameter following a Gamma distribution that involves parameter specific shrinkage and can be collected in a vector $\boldsymbol{\psi} = (\psi_1, \dots, \psi_K)'$. Overall shrinkage is governed by the global parameter λ that also follows a Gamma distribution. According to Polson and Scott (2010), this setup reflects the necessity of noise reduction by shrinking all coefficient means to zero based on the global parameter, while allowing for signals to override this effect using local scaling parameters. Hyperparameters must be set by the researcher, where standard approaches in the literature include $d_0 = d_1 = 0.01$, and $\theta = 0.1$ as default. This implies heavy overall shrinkage of the parameters stemming from the global component but provides enough flexibility to detect individual non-zero coefficients if necessary.³ Setting $\theta = 1$ would present the case leading to the Bayesian LASSO (Park and Casella, 2008). Deriving the posterior distribution of ψ_r , one finds that it follows a generalized inverse Gaussian distribution (Griffin and Brown, 2010),

$$p(\psi_r | \lambda, \beta_r) \sim \text{GIG}(\theta - 1/2, \beta_r^2, \theta \lambda^2),\tag{6}$$

with $\text{GIG}(\zeta, \chi, \varrho)$ being parameterized such that its density $f(x) \propto x^{\zeta-1} \exp\{-(\chi/x + \varrho x)/2\}$, while the conditional posterior distribution of the global parameter is a Gamma distribution with

$$p(\lambda^2 | \boldsymbol{\psi}) \sim G\left(d_0 + \theta K, d_1 + 2^{-1}\theta \sum_{r=1}^K \psi_r\right).\tag{7}$$

This setup defines the prior variance-covariance matrix on $\boldsymbol{\beta}$, denoted by $\underline{\boldsymbol{\Sigma}}$, that is updated during estimation. We set the prior coefficient vector denoted by $\underline{\boldsymbol{\beta}} = \mathbf{0}$. In the case of the NG prior, the variance-covariance matrix is given by $\text{diag}(\underline{\boldsymbol{\Sigma}}) = \boldsymbol{\psi}$. The conditional posterior for coefficients $\boldsymbol{\beta}$ is of typical form (Koop, 2003),

$$\begin{aligned}p(\boldsymbol{\beta} | \bullet) &\sim N(\bar{\boldsymbol{\beta}}, \bar{\boldsymbol{\Sigma}}), \\ \bar{\boldsymbol{\Sigma}} &= (\underline{\boldsymbol{\Sigma}}^{-1} + \mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X})^{-1}, \\ \bar{\boldsymbol{\beta}} &= \bar{\boldsymbol{\Sigma}}(\underline{\boldsymbol{\Sigma}}^{-1}\underline{\boldsymbol{\beta}} + \mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{S}(\rho)\mathbf{y}).\end{aligned}\tag{8}$$

³In addition, it is worth mentioning that θ may be also integrated out, as for instance described in Huber and Feldkircher (2017). This could easily be implemented within the given structure. However, for the sake of simplicity, we refrain from doing so.

3.2 Dirichlet–Laplace prior

Even though working well empirically, [Bhattacharya et al. \(2015\)](#) note that many aspects of global-local shrinkage priors based on scale mixtures of Gaussians are poorly understood theoretically and stress the necessity of simultaneous consideration of marginal properties of shrinkage priors and the joint distribution of obtained parameters.

As an alternative, we continue with the DL prior suggested by [Bhattacharya et al. \(2015\)](#). As opposed to the three hyperparameters to be set by the researcher in the NG case, a single hyperparameter suffices for establishing the DL setup. Similar to the NG prior, however, it is composed hierarchically of global and local shrinkage parameters, and may be structured as follows:

$$\begin{aligned}\beta_r &\sim \text{N}(0, \varphi_r \phi_r^2 \tau^2) \\ \varphi_r &\sim \text{Exp}(1/2) \\ \phi &\sim \text{Dir}(a, \dots, a) \\ \tau &\sim \text{G}(Na, 1/2).\end{aligned}\tag{9}$$

The local parameters are denoted by φ_r and assigned an exponential prior distribution. In contrast to the single global parameter λ in the case of the NG prior, the DL approach uses a vector of scales $(\phi_1 \tau, \dots, \phi_K \tau)$ to provide more flexibility regarding idiosyncratic coefficient shrinkage, where $\phi = (\phi_1, \dots, \phi_K)$ is defined to lie in the $(K - 1)$ -dimensional simplex $\mathcal{S}^{K-1} = \{\mathbf{x} = (x_1, \dots, x_K)' : x_r \geq 0, \sum_{r=1}^K x_r = 1\}$ and is assigned a Dirichlet prior distribution with hyperparameter a controlling the tightness of the prior. This quantity may again be integrated out as shown in [Bhattacharya et al. \(2015\)](#). Based on theoretical results obtained by [Bhattacharya et al. \(2015\)](#), standard deterministic choices include $a = 1/K$ (the default setting used for the simulation study), but different choices such as $a = 1/2$ are justifiable theoretically. Notice that this allows the DL prior setup to be dependent on the dimensionality of the underlying model on theoretical grounds, different to both the SSVS and NG prior.

In the case of the DL prior, [Bhattacharya et al. \(2015\)](#) show that φ_r can be sampled efficiently involving independent inverse Gaussian distributions. This is done by obtaining $\tilde{\varphi}_r | \phi, \beta$ from a reparameterization of the generalized inverse Gaussian distribution with $\mu_r = \phi_r \tau / |\beta_r|$,

$$p(\tilde{\varphi}_r | \phi, \beta) \sim \text{GIG}(-1/2, 1, \mu_r^{-2}),\tag{10}$$

where we subsequently set $\varphi_r = 1/\tilde{\varphi}_r$ to obtain draws from the conditional posterior distribution of φ_r . The global shrinkage component τ is again sampled from a generalized inverse Gaussian distribution

$$p(\tau | \phi, \beta) \sim \text{GIG}\left(1 - K, 2 \sum_{r=1}^K |\beta_r| / \phi_r, 1\right).\tag{11}$$

In order to sample $\phi | \beta$ we rely on Theorem 2.1 in [Bhattacharya et al. \(2015\)](#) and draw auxiliary variables T_1, \dots, T_K independently, with $T_r \sim \text{GIG}(a - 1, 2|\beta_r|, 1)$. Afterwards, we set $\phi_r = T_r / \sum_{j=1}^K T_j$. [Bhattacharya et al. \(2015\)](#) indicate this step as an important feature of their setup, as it significantly accelerates mixing and convergence. Consequently, the structure of the prior variance-covariance matrix $\underline{\Sigma}$ is updated, where $\text{diag}(\underline{\Sigma}) = (\varphi_1 \phi_1^2 \tau^2, \dots, \varphi_K \phi_K^2 \tau^2)'$. The coefficient vector is then sampled analogously to the NG prior, based on [Eq. \(8\)](#).

3.3 Priors for remaining parameters

Without loss of generality, we assume a homoskedastic error variance $\Omega = \sigma^2 \mathbf{I}_N$, where \mathbf{I}_N is an $N \times N$ -dimensional identity matrix.⁴ To complete the prior setup, we have to elicit prior distributions for σ^2 and the spatial dependence parameter ρ . Here we use standard specifications in the literature (LeSage and Pace, 2009). Specifically, we impose an inverse Gamma prior on the variance of the error term, $\sigma^2 \sim G^{-1}(\underline{a}, \underline{b})$ with scalar hyperparameters \underline{a} and \underline{b} that we set equal to 0.01, rendering them rather uninformative. Conditioning on all other parameters of the model and the data we obtain a conditional posterior density for the error variances,

$$\begin{aligned} p(\sigma^2|\bullet) &\sim G^{-1}(\bar{a}, \bar{b}), \\ \bar{a} &= \underline{a} + N/2, \\ \bar{b} &= \underline{b} + \boldsymbol{\varepsilon}'\boldsymbol{\varepsilon}/2. \end{aligned} \tag{12}$$

The quantities for $\boldsymbol{\beta}$ and σ^2 are standard and can be sampled efficiently using Gibbs sampling (Koop 2003). For ρ , we follow LeSage and Pace (2007) and LeSage and Pace (2009), by eliciting a normal prior $\rho \sim N(0, c)$, where c may be chosen dependent on the prior belief of the researcher regarding the strength of spatial association in the data. For the simulation study, we again use a rather uninformative specification and choose $c = 10$. Since the posterior distribution of ρ conditional on all other quantities of the model is in general not of well-known form,

$$p(\rho|\bullet) \propto \exp(-\boldsymbol{\varepsilon}'\boldsymbol{\Omega}^{-1}\boldsymbol{\varepsilon}/2)p(\rho), \tag{13}$$

it is sampled by a Metropolis-within-Gibbs step. We follow the standard approach by proposing a new value from a Normal distribution, $\rho^* \sim N(\rho_{t-1}, \zeta)$, where the subscript $t - 1$ denotes the value of the parameter from the previous iteration of the sampling algorithm and ζ is a tuning parameter. The acceptance probability of the proposal is calculated using

$$\min \left[1, \frac{p(\rho^*|\bullet)}{p(\rho_{t-1}|\bullet)} \right].$$

If the proposal is accepted, we set $\rho_t = \rho^*$. Otherwise, the proposal is rejected, and we retain the value from the previous draw. The process of generating proposals for ρ is tuned during half of the burn-in phase of the algorithm by incrementally increasing or decreasing the variance of the proposal distribution ζ to yield an acceptance rate for the parameter between 0.2 and 0.4. An overview of the algorithm employed can be found in Appendix B.

4 Simulation study

In this section, we evaluate the empirical properties and merits of the proposed shrinkage priors. Estimates are obtained using the Normal-Gamma (NG) and Dirichlet-Laplace (DL) shrinkage priors sketched above. As a benchmark specification, we compare the results of both setups with a stochastic search variable selection (SSVS) prior put forward by George and McCulloch (1993) and applied to spatial autoregressive models by Piribauer (2016) and Piribauer and Crespo Cuaresma (2016). The SSVS prior as a means to introducing shrinkage on the slope coefficients mimics Bayesian model-averaging frameworks.

⁴The generic notation in Eq. (1), however, allows for various specifications of the variance-covariance matrix Ω , for instance reflecting heteroskedastic error terms (see, for instance, LeSage, 1997, in a spatial context).

Details on the SSVS setup along with the concurrent Markov-chain Monte Carlo (MCMC) sampling algorithm are given in Appendix A. To further illustrate and underline the necessity of applying shrinkage in cases where the coefficient vector is sparse, we also provide results for the case where a rather uninformative prior variance-covariance matrix $\underline{\Sigma} = 1000\mathbf{I}_N$ is used. This approach represents the basic setup given in LeSage and Pace (2007), and is labeled *None*, referring to the fact that no shrinkage is applied. For each specification employed, we use 2,000 iterations, discarding the initial 1,000 draws as burn-in. Inference is then based on the $T = 1,000$ posterior draws for all parameters, where point estimates are calculated using the median of the respective sample.

4.1 Data Generating Process

Simulating a synthetic data set for different data generating processes requires decisions on the number of observations N and explanatory variables K . Moreover, we have to set parameter values for the coefficients $\tilde{\boldsymbol{\beta}}$, the variance of the error terms $\tilde{\sigma}^2$ and the spatial autoregressive parameter $\tilde{\rho}$. We consider varying degrees of sparsity with respect to the coefficient vector. The final data generating process, reflecting the basic modeling approach as in Eq. (1), has the form

$$\tilde{\mathbf{y}} = \exp(-\tilde{\rho}\tilde{\mathbf{W}})(\tilde{\mathbf{Z}}\tilde{\boldsymbol{\beta}} + \boldsymbol{\epsilon}), \quad \tilde{\boldsymbol{\epsilon}} \sim \text{N}(0, \tilde{\sigma}^2\mathbf{I}_N), \quad (14)$$

where the first column in \mathbf{Z} contains an intercept. Robustness for different combinations of parameters and generated variables is achieved via repeating each of the basic simulations sketched below for 100 times and relying on stochastically setting most of the quantities involved. In particular, we set the required parameters as follows:

- (i) The values for each of the $k = 1, \dots, K-1$ (non-constant) explanatory variables $\mathbf{z}_k = [z_{k1}, \dots, z_{kN}]'$ for each observation $i = 1, \dots, N$ come from a standard normal distribution, $z_{ki} \sim \text{N}(0, 1)$ and are collected in the matrix $\mathbf{Z} = (\boldsymbol{\iota}_N, \mathbf{z}_1, \dots, \mathbf{z}_{K-1})$ where $\boldsymbol{\iota}_N$ is an N -dimensional vector of ones. For all simulations, we use a sample size of $N = 100$ and consider $K \in \{50, 100, 150, 200\}$.
- (ii) For the longitude and latitude coordinates of the observations we randomly generate points within a unit square and construct a row-stochastic matrix \mathbf{W} based on a five-nearest neighborhood specification.⁵ It is worth mentioning that the specification of the number of neighbors appeared to have no discernible effect on the results obtained during simulations.
- (iii) Sparsity of the coefficient vector is governed by setting a maximum number of non-zero parameters. We evaluate two different scenarios, where the number of $\tilde{\beta}_k \neq 0$ does not exceed $q \in \{10, 20\}$. As we also intend to test adaptiveness of shrinkage, half of the non-zero parameters are generated to be close to zero; the other half is allowed to exhibit greater variation, while we always include a non-zero effect on the intercept. For the slope coefficients corresponding to the non-constant covariates in \mathbf{Z} , this is achieved by simulating $(K-1)$ true coefficient values $\tilde{\beta}_{1:(q/2)} \sim \text{N}(0, 1)$ and $\tilde{\beta}_{(q/2+1):q} \sim \text{N}(0, 5)$, where the parameter for the intercept always comes from the latter distribution. The remaining coefficients are set to zero in a deterministic fashion, $\tilde{\beta}_{(q+1):K} = 0$.
- (iv) The homoskedastic variance of the error terms is deterministically set to $\sigma^2 = 1$.
- (v) The spatial dependence parameter is simulated using a zero-mean Gaussian distribution, $\tilde{\rho} \sim \text{N}(0, 3)$.

⁵Along the lines of Strauss et al. (2017) we do not use a truncated Taylor series expansion of the matrix exponential for calculations in this article, but an efficient implementation of this operation in R.




Fig. 1: Diagnostic plots for a true coefficient equal to zero ($N = 100$, $K = 50$ and $q = 10$).

Notes: The panel above shows trace plots for the specifications under scrutiny (horizontal axis: number of iteration; vertical axis: parameter value), while the lower panel depicts the posterior distribution of the parameter (horizontal axis: parameter value; vertical axis: density). The dashed red line indicates the true parameter value, the blue line is the median of the obtained posterior distribution.

4.2 Diagnostics

We briefly consider diagnostics regarding obtained draws from the posterior distributions by evaluating trace plots and estimated densities. This allows us to present both the obtained marginal posterior densities for the respective parameters, and to discuss mixing and convergence properties. Diagnostic plots of non-zero coefficients typically mirror the well-known picture resulting from standard applications without introducing shrinkage. For illustrative purposes, we thus pick examples with true coefficients being equal to zero. A low-dimensional problem is given by choosing $q = 10$ and $K = 50$, while the case of $K = 150$ serves as high-dimensional scenario where applying shrinkage is required to find meaningful results on an acceptable level of precision.

The first example, for $K = 50$ exogenous covariates, is presented in Fig. 1. The upper panel shows trace plots for all prior specifications, while the lower depicts the marginal posterior density of the parameter. The dashed red line indicates the true parameter value, the blue line is the median of the obtained posterior distribution. Even in this case, where ten out of the 50 slope coefficients in the parameter vector β are different from zero, we find that the standard specification without shrinkage (*None*) of the parameter space results in a comparatively high variance of parameter estimates, observable in Fig. 1. Considering the three shrinkage approaches, we find that the DL and NG priors concentrate more point mass on zero by design, while the SSVS prior shows considerable variation in the interval around zero.

A more interesting case is given in Fig. 2. In this setup we include 150 covariates, where only ten exhibit non-zero values by construction. As is evident considering *None*, the variance around parameter estimates is huge when the number of parameters to estimate exceeds the number of observations. Since the semi-automatic setup of the SSVS prior set forth in Appendix A relies on these quantities to scale prior hyperparameters, this is influential regarding resulting parameter estimates for the SSVS approach.

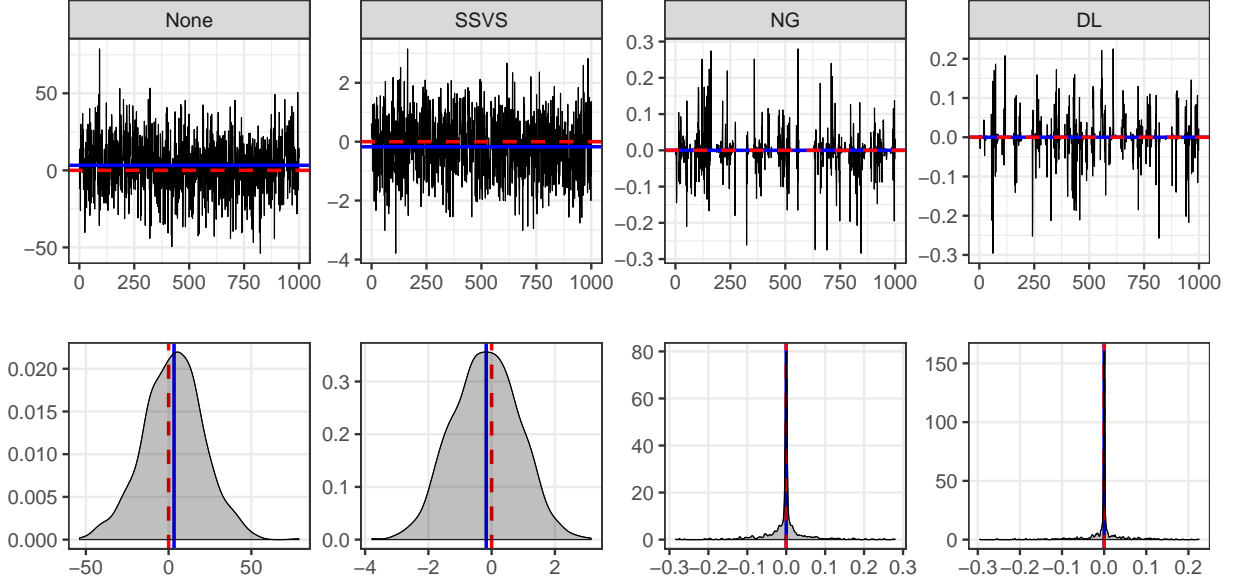


Fig. 2: Diagnostic plots for all prior setups of a true coefficient equal to zero ($K = 150$ and $q = 10$).

Notes: The panel above shows trace plots for the specifications under scrutiny (horizontal axis: number of iteration; vertical axis: parameter value), while the lower panel depicts the posterior distribution of the parameter (horizontal axis: parameter value; vertical axis: density). The dashed red line indicates the true parameter value, the blue line is the median of the obtained posterior distribution.

Notice that the obtained scaling factor renders the variance of the spike component comparatively large due to the specific setup involved, and SSVS does not provide enough shrinkage, as shown in Fig. 2.

This issue is not observable in the case of the DL and NG priors, providing evidence for their excellent adaptive shrinkage properties in high-dimensional settings. For the NG prior, it is worth mentioning that the point mass placed on zero is comparable to the case for $K = 50$. By the fact that the single hyperparameter of the DL prior is set to $1/K$ in the default case, higher dimensional problems result in even stronger shrinkage towards zero for true zero coefficients. Evidence of the increased tightness of the prior can be observed with regard to the density scales in Fig. 2. Moreover, note that we also observe that the shrinkage priors are able to recover coefficients close to zero equally well.

Following this brief description of diagnostic plots, we again use the scenarios above as examples and consider all of the resulting parameter estimates jointly. Plots of the true regression coefficients against their posterior medians are depicted in Fig. 3 and Fig. 4. The black line indicates the 45 degree line, reflecting the objective of perfectly estimated parameters. In the low-dimensional example given in Fig. 3, we find that point estimates based on the median of obtained posterior distributions for the parameters are mostly correct also in the case where no shrinkage is applied. However, note that these estimates are rather imprecise, which renders them suboptimal in terms of significance for interpretation, and also for applications involving predictions. The three shrinkage approaches closely mirror each other, even though it is worth noting that the NG and DL prior allow for even more precise estimates in terms of the variance.

A completely different case emerges for 150 covariates. Fig. 4 showcases severe problems regarding the standard prior specification without shrinkage (*None*). Even though the SSVS prior performs slightly better than this approach, it is evident that the default setup regarding the scaling of the hyperparameters




Fig. 3: Comparison of posterior estimates to true parameter values for $K = 50$ and $q = 10$.

Notes: The panel above shows the correlation between true parameter values and estimated quantities based on the mean of the posterior draws for scenario 1. The black line indicates the 45 degree line.




Fig. 4: Comparison of posterior estimates to true parameter values for $K = 150$ and $q = 10$.

Notes: The panel above shows the correlation between true parameter values and estimated quantities based on the mean of the posterior draws. The black line indicates the 45 degree line.

is suboptimal in high-dimensions. In particular, and as we will discuss below, the SSVS prior actually tracks non-zero coefficients quite well, but imprecisely. Issues arise mainly from its disability to capture zero coefficients adequately in the given amount of iterations of the MCMC algorithm, pointing towards mixing problems and suboptimal convergence as stated in [Bhattacharya et al. \(2015\)](#). In this respect, both the NG and DL prior are superior. The extraordinary empirical properties of these priors are evident in [Fig. 4](#), where the majority of the coefficients is recovered almost perfectly. Note that the DL prior provides slightly tighter shrinkage of true zero coefficients.

4.3 Evaluation of performance

Turning to an overview of all estimated parameters of the models, with respect to a total of eight different scenarios given by combinations of $K = \{50, 100, 150, 200\}$ and $q = \{10, 20\}$, we present root mean squared error (RMSE) measures.⁶ The specific choice of the number of included covariates and number

⁶We adopt the standard root mean squared error measure for the posterior median $\hat{\beta}$, $\text{RMSE}(\hat{\beta}) = (\sum_{k=1}^K (\hat{\beta}_k - \tilde{\beta}_k)^2 / K)^{\frac{1}{2}}$ and also calculate the corresponding quantity with respect to each draw from the MCMC algorithm, $\text{RMSE}_{\text{dr}}(\hat{\beta}) = \sum_{k=1}^K (\sum_{t=1}^T (\hat{\beta}_{k,t} - \tilde{\beta}_k)^2 / T)^{\frac{1}{2}} / K$. This serves as a means to illustrate differing degrees of precision of the estimators.

Table 1: Root mean squared error measures for point estimates.

	RMSE	$q = 10$				$q = 20$			
		β	σ^2	ρ	Time	β	σ^2	ρ	Time
$K = 50$	None	0.0184	0.2319	0.0044	1.02	0.0180	0.3072	0.0028	1.00
	SSVS	0.0078	0.0150	0.0032	1.00	0.0107	0.0189	0.0021	1.00
	NG	0.0090	0.0155	0.0034	1.05	0.0111	0.0168	0.0022	1.06
	DL	0.0089	0.0185	0.0035	1.01	0.0120	0.0311	0.0021	1.02
$K = 100$	None	0.4589	> 10	0.1606	1.02	0.4670	> 10	0.1191	1.04
	SSVS	0.0118	0.0398	0.0038	1.09	0.0130	0.0317	0.0025	1.06
	NG	0.0080	0.0297	0.0033	1.06	0.0101	0.0296	0.0021	1.02
	DL	0.0073	0.0182	0.0032	1.00	0.0101	0.0316	0.0023	1.00
$K = 150$	None	0.3665	> 10	0.2914	1.00	0.4847	> 10	0.2438	1.00
	SSVS	0.2578	0.0993	0.0304	1.11	0.1882	0.0973	0.0265	1.09
	NG	0.0075	0.0498	0.0035	1.11	0.0088	0.0500	0.0026	1.09
	DL	0.0064	0.0226	0.0038	1.06	0.0080	0.0304	0.0023	1.03
$K = 200$	None	0.2885	> 10	0.2625	1.00	0.3720	> 10	0.2522	1.00
	SSVS	0.1353	0.0958	0.0503	1.21	0.1354	0.1103	0.0367	1.18
	NG	0.0075	0.0789	0.0039	1.12	0.0087	0.0799	0.0027	1.09
	DL	0.0055	0.0268	0.0032	1.02	0.0074	0.0262	0.0025	1.00

Notes: K denotes the number of covariates, q is the number of non-zero parameters. Time refers to the relative time with respect to the fastest approach for the given scenario. RMSE is the root mean squared error.

of non-zero coefficients implies that we have degrees of sparsity ranging from 5 to 40 percent for the parameter vector. Average results for 100 iterations are presented in [Tab. 1](#) and [Tab. 2](#).

The results for the standard RMSE measure closely resemble the notions obtained in the discussion above. In particular, we find that the SSVS prior performs slightly better than other shrinkage approaches in low dimensional settings where $K < N$, both in terms of estimation errors for parameters in β , σ^2 , the spatial dependence parameter ρ and also time elapsed for cycling through the MCMC algorithm. However, both the NG and DL prior show promising values regarding RMSEs and can be considered as feasible alternatives. For the case of $K = N$, we find a different picture. The approach without shrinkage (*None*) performs rather poorly, evidenced by large errors in terms of all parameters. In particular, based on poor estimates for the coefficients in β , this results in the inability to correctly estimate σ^2 . We find particularly good performance measures for the DL prior with the NG prior being close second. This finding is particularly pronounced in the denser case with respect to the coefficient vector ($q = 20$), where the NG prior outperforms the DL prior in terms of precision regarding the error variances and the spatial dependence parameter.

Turning to the higher dimensional cases where K exceeds N it is worth noting that SSVS runs into similar problems than the standard approach, where poor parameter estimates result in huge values for the error variances. Since this problem only occurred in a minor subset of the 100 iterations per scenario, we chose to exclude these erroneous simulations. In this setup, the approach without shrinkage again appeared to perform poorly, as already observed in the case with $K = 100$. Second, due to the notion that SSVS typically shows suboptimal mixing and convergence properties in high dimensional settings – stemming from the necessity of stochastic search over an enormous parameter space – it appears that the default setting of 2,000 iterations for the MCMC algorithm is not sufficient to obtain reasonable posterior distributions for the parameters. Finally, we find that the DL prior performs best regarding coefficients and error variances – compared to the NG prior, which is only superior in terms of estimating the spatial dependence parameter ρ . Note that estimates in the case of more dense coefficient vectors ($q = 20$)

Table 2: Root mean squared error measures for density estimates.

	RMSE _{dr}	$q = 10$				$q = 20$			
		β	σ^2	ρ	Time	β	σ^2	ρ	Time
$K = 50$	None	0.0250	0.0371	0.0072	1.02	0.0248	0.0363	0.0043	1.00
	SSVS	0.0096	0.0273	0.0051	1.00	0.0127	0.0318	0.0034	1.00
	NG	0.0122	0.0276	0.0053	1.05	0.0149	0.0300	0.0034	1.06
	DL	0.0130	0.0357	0.0058	1.01	0.0174	0.0513	0.0039	1.02
$K = 100$	None	0.6730	> 10	0.2057	1.02	0.6795	> 10	0.1482	1.04
	SSVS	0.0171	0.0480	0.0058	1.09	0.0185	0.0452	0.0040	1.06
	NG	0.0113	0.0381	0.0052	1.06	0.0136	0.0398	0.0034	1.02
	DL	0.0118	0.0367	0.0056	1.00	0.0154	0.0569	0.0041	1.00
$K = 150$	None	2.1817	> 10	0.3142	1.00	2.1749	> 10	0.2807	1.00
	SSVS	0.1812	0.1034	0.0314	1.11	0.1376	0.1141	0.0276	1.09
	NG	0.0106	0.0549	0.0053	1.11	0.0122	0.0553	0.0036	1.09
	DL	0.0111	0.0396	0.0061	1.06	0.0138	0.0596	0.0042	1.03
$K = 200$	None	2.4875	> 10	0.3122	1.00	2.3582	> 10	0.3091	1.00
	SSVS	0.0954	0.1424	0.0516	1.21	0.1106	0.1806	0.0377	1.18
	NG	0.0103	0.0812	0.0051	1.12	0.0117	0.0822	0.0036	1.09
	DL	0.0111	0.0422	0.0053	1.02	0.0138	0.0631	0.0044	1.00

Notes: K denotes the number of covariates, q is the number of non-zero parameters. Time refers to the relative time with respect to the fastest approach for the given scenario. RMSE_{dr} is the root mean squared error based on all MCMC draws providing insights in the performance in terms of parameter precision.

are typically slightly worse, due to the necessity of estimating more none-zero parameters with the same number of observations. A similar picture is present in the case of $K \gg N$. Mirroring the results already obtained in the previous scenario, the DL prior appears to be the fastest of the shrinkage priors. Interestingly, the superiority of the DL prior regarding estimates of the error variances gets even larger when compared to the worse estimates resulting from the NG prior.

The main results discussed above also hold in the case of the adapted RMSE focusing on density predictions in Tab. 2. Recall that we use this measure for the purpose of gaining insight into the precision of the obtained parameter estimates. Interestingly, even though being inferior to the DL prior in terms of point estimates, we find that the NG prior produces more precise estimates around the true values of the parameters in the coefficient vector β . However, this finding does not carry over to estimates of the error variances σ^2 , where the DL prior outperforms the NG approach. The spatial dependence parameter is estimated precisely in all cases regarding sparsity and number of covariates within the NG and DL shrinkage prior framework, and in scenarios where the SSVS prior and the standard approach still perform well.

5 Empirical illustration

In this section we aim at illustrating the performance of the proposed model specification using real data on pan-European regional economic growth and its empirical determinants. Specifically, we consider a cross-regional spatial Durbin model specification (see, for example, LeSage and Pace 2009) which also allows for spatially lagged explanatory variables as potential covariates. The specification used can be written as

$$S(\rho)y = \iota_N \alpha + X\beta + WX\vartheta + u, \quad u \sim N(\mathbf{0}, \sigma^2 I_N) \quad (15)$$

where \mathbf{y} denotes an N -dimensional column vector of regional economic growth rates of per capita gross value added and the $N \times 1$ error term \mathbf{u} is defined as before as iid normal. The $N \times K$ matrix \mathbf{X} contains the set of potential predictors. \mathbf{W} is an $N \times N$ row-stochastic spatial weight matrix as defined before. $S(\rho)$ denotes a matrix exponential spatial filter with corresponding scalar parameter ρ as defined before. \mathbf{WX} is the spatial lag of the explanatory variables and explicitly incorporates spatially lagged information of \mathbf{X} from neighboring regions, resulting in the spatial Durbin model specification mentioned above. In this empirical illustration we follow the typical structure of (spatial) growth regressions by assuming that information in matrix \mathbf{X} is measured at the beginning of the sample period (which is the year 2000).

5.1 Regions, spatial weights and data

For the empirical illustration we use data on regional economic growth on a sample of 273 European NUTS-2 regions of 28 European countries. The dependent variable in the regression framework is the average annual growth rate of per capita gross value added in the period 2001 – 2010. A detailed list of the regions used in the application is given in [Tab. C.2](#) in [Appendix C](#). [Table C.1](#) provides detailed information on the set of potential covariates in the matrix exponential spatial growth specification.

The set of predictors is in line with recent empirical applications on regional economic growth. Specifically, the matrix of explanatory variables contains information on the initial level of economic growth rates, human capital endowments, proxies for knowledge capital stocks, regional population structure, infrastructure, the region-specific industry mix, and other socio-economic quantities. It is worth noting that the spatial lag of the (non-constant) explanatory variables in the spatial Durbin framework sketched above results in doubling the set of potential covariates depicted in [Tab. C.1](#). For the specification of the spatial weight matrix \mathbf{W} we used a 10-nearest row-stochastic specification.⁷

5.2 Empirical results

In this subsection we present the results of the stochastic search variable selection (SSVS) prior, along with Normal-Gamma (NG), and Dirichlet-Laplace (DL) shrinkage setups including alternative prior hyperparameter specifications. Specifically, for the Normal-Gamma shrinkage prior a natural candidate is the Bayesian LASSO ([Park and Casella, 2008](#)), achieved by setting the prior hyperparameter $\theta = 1$. Alternatively, we also consider $\theta = 0.1$, which refers to the generalized version characterized by heavier overall shrinkage. For the Dirichlet-Laplace (DL) shrinkage prior, we consider $a = 1/2$, which presents a usual benchmark in empirical research, while $a = 1/K$ is the default prior specification based on theoretical considerations (see [Bhattacharya et al., 2015](#)).

[Table 3](#) depicts estimation results for the competing prior setups under scrutiny. Note that only covariates where the corresponding slope coefficient is statistically significant from zero based on the lower 10 percent and upper 90 posterior interval in at least one of the candidate specifications are reported. For each prior setup, [Tab. 3](#) reports the posterior means (labeled Mean) and corresponding posterior standard deviations (labeled SD) for the parameters under scrutiny. Both quantities are based on 3,000 retained draws of the MCMC algorithms described in [Appendix B](#).⁸

Highly significant estimates for the spatial parameter ρ result in all specifications considered, ranging from -0.811 to -0.972 . Using the correspondence between the spatial parameter (ξ) in conventional spatial autoregressive frameworks and its matrix exponential spatial counterpart (ρ) stated by [LeSage](#)

⁷Robustness checks using alternative numbers of nearest neighbors appeared to exert negligible effects on the results.

⁸Note that for all alternative specifications we used a number of 12,000 iterations with the first 3,000 serving as burn-ins. To reduce the autocorrelation of the draws for ρ we have used *thinning* by considering only every third draw (see, for example, [Koop 2003](#)), resulting in a total of 3,000 draws for posterior inference.

and Pace (2007), we find an implied spatial autoregressive parameter ξ between 0.56 and 0.62. This degree of spatial dependence resembles the findings in recent studies on regional spatial economic growth in Europe (see, for example, Crespo Cuaresma et al., 2014; LeSage and Fischer, 2008).

Turning attention to the (non-constant) potential growth determinants, Tab. 3 shows four slope coefficients that are statistically significant in all prior setups. These variables are the initial level of income, the old-age dependency ratio, as well as both the share of low-educated employment and its corresponding spatial lag. Overall, Tab. 3 shows rather similar results for both the magnitudes of the estimated effects as well as their significance. However, a notable exception is the Bayesian LASSO setup $NG(\theta = 1)$, which highlights a markedly higher amount of significant slope parameters as compared to the competing specifications.

As expected, the initial income variable appears to negatively affect regional economic growth rates in all specifications, providing evidence for conditional convergence among the regions in the sample. However, the SSVS setting also shows significant posterior mean of the spatially lagged initial income variable, pointing towards the existence of positive growth spillovers emanating from the level of income of neighboring regions. This means that regions seem to benefit from the spatial proximity of rich regions in terms of accelerated growth rates. The old-age dependency ratio (measured in terms of the ratio of the number of people aged 65 and over to the working age population) appears to exhibit a negative impact on regional economic growth rates. Except for the Bayesian LASSO specification ($NG(\theta = 1)$), the estimated coefficients are rather similar.

Interestingly, the results presented in Tab. 3 corroborate the findings of previous studies (see, for example, Piribauer and Fischer 2015, or Crespo Cuaresma et al. 2014) by detecting the share of low educated working age population (Lower education workers) as being more robustly correlated to regional economic growth as compared to a measure of tertiary education attainment (higher education workers). As expected, lower education workers appear to exhibit a negative impact on regional growth in all specifications considered. However, it is worth noting that the spatial lag of this variable appears to exhibit a positive impact, indicating that positive effects on income growth to neighboring regions. Work by Olejnik (2008), for example, argue that an increase in the lower educated labor force might result in positive growth spillovers, by assuming that such an increase might be primarily due to migration of workers between regions.

6 Concluding remarks

Dealing with model uncertainty in spatial autoregressive model specifications has been subject to numerous studies in general, especially in the empirical economic growth literature. However, spatial econometric applications typically rely on Bayesian model-averaging techniques, which suffer from severe drawbacks both in terms of computational time and possible extensions to more flexible model specifications. In spatial autoregressive models, the computational burden emanates from the calculation of marginal likelihoods, where no closed form solutions are available. Recent contributions to the literature as a means to alleviating the computational difficulties include a variant of the conventional stochastic search variable selection prior discussed in Piribauer and Crespo Cuaresma (2016). However, shortcomings of this approach include slow mixing and convergence properties in the presence of a large number of available covariates and difficulties in choosing prior hyperparameters.

This paper aims at generalizing two absolutely continuous hierarchical shrinkage priors – the Normal-Gamma and the Dirichlet-Laplace shrinkage prior – to the matrix exponential spatial specification in order to alleviate the above-mentioned drawbacks of both Bayesian model averaging and standard stochastic search variable selection priors. The proposed frameworks allow for flexible and adaptive, but also computationally efficient stochastic variable selection, where extensions to basic spatial model

Table 3: Estimation results.

Variable	SSVS		NG ($\theta = 1$)		NG ($\theta = 0.1$)		DL ($a = 1/2$)		DL ($a = 1/K$)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Initial income*	-0.857	0.195	-0.702	0.194	-0.734	0.241	-0.669	0.225	-0.702	0.209
Population density	-0.036	0.062	-0.205	0.107	-0.050	0.074	-0.036	0.068	-0.031	0.063
Child dependency ratio	-0.017	0.025	-0.057	0.028	-0.018	0.026	-0.011	0.021	-0.009	0.019
Old-age dependency ratio*	-0.038	0.022	-0.055	0.020	-0.036	0.022	-0.028	0.022	-0.027	0.022
Lower education workers*	-0.060	0.009	-0.060	0.011	-0.056	0.012	-0.048	0.016	-0.051	0.014
<i>W</i> Initial income	0.569	0.388	0.173	0.248	0.251	0.294	0.140	0.242	0.178	0.270
<i>W</i> Child dependency ratio	0.021	0.033	0.080	0.054	0.020	0.035	0.015	0.030	0.012	0.027
<i>W</i> Agriculture employment	0.007	0.012	0.033	0.025	0.002	0.010	0.002	0.011	0.001	0.009
<i>W</i> Construction employment	0.019	0.037	0.123	0.088	0.013	0.037	0.018	0.042	0.014	0.037
<i>W</i> Technology resources	0.004	0.012	0.054	0.035	0.003	0.011	0.002	0.013	0.002	0.010
<i>W</i> Higher education workers	-0.003	0.012	-0.050	0.032	-0.002	0.010	-0.003	0.013	-0.001	0.009
<i>W</i> Lower education workers*	0.048	0.013	0.055	0.015	0.043	0.015	0.030	0.020	0.033	0.019
σ^2	0.642	0.057	0.637	0.060	0.657	0.081	0.671	0.065	0.667	0.064
ρ	-0.972	0.151	-0.811	0.146	-0.925	0.138	-0.848	0.134	-0.874	0.146

Notes: The table contains a subset of all variables for which coefficients are significant in at least one specification (*significant in all specifications). Estimates depicted in bold are statistically different from zero based on the 80 percent posterior credible interval.

specification can be easily implemented in Markov chain Monte Carlo sampling algorithms. For illustrative purposes, and to evaluate prior-specific properties in the presence of spatial dependence, the paper carries out an extensive simulation study using synthetic data sets. An empirical illustration is given by a study on economic growth of European regions.

Our results indicate that standard stochastic search variable selection priors work particularly well in relatively low-dimensional settings. However, severe problems occur in high-dimensional settings, where the number of potential covariates relative to the number of available observations becomes large. The proposed global-local shrinkage priors provide the required adaptiveness of shrinkage in a flexible and computationally efficient way. The Normal-Gamma shrinkage prior excels in terms of precision of estimates, while the Dirichlet-Laplace shrinkage prior slightly outperforms the former in terms of point estimations of parameters. Both proposed shrinkage priors can be considered valuable tools as a means to incorporate model uncertainty in spatial autoregressive frameworks, regardless of dimensionality of the problem at hand.

References

- BHATTACHARYA, A., D. PATI, N. S. PILLAI, AND D. B. DUNSON (2015): “Dirichlet-Laplace priors for optimal shrinkage,” *Journal of the American Statistical Association*, 110(512), 1479–1490.
- BITTO, A., AND S. FRÜHWIRTH-SCHNATTER (2016): “Achieving shrinkage in a time-varying parameter model framework,” Discussion paper, ArXiv:1611.01310.
- CORNWALL, G. J., AND O. PARENT (2017): “Embracing heterogeneity: The spatial autoregressive mixture model,” *Regional Science and Urban Economics*, 64, 148–161.
- CRESPO CUARESMA, J., G. DOPPELHOFER, AND M. FELDKIRCHER (2014): “The determinants of economic growth in european regions,” *Regional Studies*, 48(1), 44–67.
- CRESPO CUARESMA, J., G. DOPPELHOFER, F. HUBER, AND P. PIRIBAUER (2018): “Human capital accumulation and long-term income growth projections for European regions,” *Journal of Regional Science*.
- CRESPO CUARESMA, J., AND M. FELDKIRCHER (2013): “Spatial filtering, model uncertainty and the speed of income convergence in Europe,” *Journal of Applied Econometrics*, 28(4), 720–741.

- DEBARSY, N., F. JIN, AND L. FEI LEE (2015): “Large sample properties of the matrix exponential spatial specification with an application to FDI,” *Journal of Econometrics*, 188(1), 1 – 21.
- FELDKIRCHER, M., F. HUBER, AND G. KASTNER (2017): “Sophisticated and small versus simple and sizeable: When does it pay off to introduce drifting coefficients in Bayesian VARs?,” Discussion paper, ArXiv:1711.00564v2.
- GEORGE, E. I., AND R. E. MCCULLOCH (1993): “Variable selection via Gibbs sampling,” *Journal of the American Statistical Association*, 88(423), 881–889.
- (1997): “Approaches for Bayesian variable selection,” *Statistica Sinica*, 7(2), 339–373.
- GRIFFIN, J. E., AND P. J. BROWN (2010): “Inference with normal-gamma prior distributions in regression problems,” *Bayesian Analysis*, 5(1), 171–188.
- (2017): “Hierarchical shrinkage priors for regression models,” *Bayesian Analysis*, 12(1), 135–159.
- HALLECK VEGA, S., AND J. P. ELHORST (2015): “The SLX model,” *Journal of Regional Science*, 55(3), 339–363.
- HÖRMANN, W., AND J. LEYDOLD (2014): “Generating generalized inverse Gaussian random variates,” *Statistics and Computing*, 24(4), 547–557.
- HUBER, F., AND M. FELDKIRCHER (2017): “Adaptive shrinkage in Bayesian vector autoregressive models,” *Journal of Business & Economic Statistics*, forthcoming.
- KASTNER, G. (2016): “Sparse Bayesian time-varying covariance estimation in many dimensions,” Discussion paper, ArXiv:1608.08468, forthcoming in the *Journal of Econometrics*.
- KOOP, G. (2003): *Bayesian econometrics*. John Wiley & Sons, Chichester, UK.
- LAURINI, M. P. (2017): “A spatial error model with continuous random effects and an application to growth convergence,” *Journal of Geographical Systems*, 19(4), 371–398.
- LESAGE, J. P. (1997): “Bayesian estimation of spatial autoregressive models,” *International Regional Science Review*, 20(1-2), 113–129.
- LESAGE, J. P., AND Y.-Y. CHIH (2017): “A matrix exponential spatial panel model with heterogeneous coefficients,” *Geographical Analysis*, forthcoming.
- LESAGE, J. P., AND M. M. FISCHER (2008): “Spatial growth regressions: Model specification, estimation and interpretation,” *Spatial Economic Analysis*, 3(3), 275–304.
- LESAGE, J. P., AND R. K. PACE (2007): “A matrix exponential spatial specification,” *Journal of Econometrics*, 140(1), 190–214.
- (2009): *Introduction to spatial econometrics*. Taylor & Francis Group, LLC, Boca Raton, FL.
- LESAGE, J. P., AND O. PARENT (2007): “Bayesian model averaging for spatial econometric models,” *Geographical Analysis*, 39(3), 241–267.
- OLEJNIK, A. (2008): “Using the spatial autoregressively distributed lag model in assessing the regional convergence of per-capita income in the EU25,” *Papers in Regional Science*, 87(3), 371–384.
- PARK, T., AND G. CASELLA (2008): “The Bayesian Lasso,” *Journal of the American Statistical Association*, 103(482), 681–686.
- PIRIBAUER, P. (2016): “Heterogeneity in spatial growth clusters,” *Empirical Economics*, 51(2), 659–680.
- PIRIBAUER, P., AND J. CRESPO CUARESMA (2016): “Bayesian variable selection in spatial autoregressive models,” *Spatial Economic Analysis*, 11(4), 457–479.
- PIRIBAUER, P., AND M. M. FISCHER (2015): “Model uncertainty in matrix exponential spatial growth regression models,” *Geographical Analysis*, 47(3), 240–261.
- POLSON, N. G., AND J. G. SCOTT (2010): “Shrink globally, act locally: sparse Bayesian regularization and prediction,” *Bayesian Statistics*, 9, 501–538.
- STEEL, M. F. (2017): “Model averaging and its use in economics,” *arXiv preprint arXiv:1709.08221*.
- STRAUSS, M. E., M. MEZZETTI, AND S. LEORATO (2017): “Is a matrix exponential specification suitable for the modeling of spatial correlation structures?,” *Spatial Statistics*, 20, 221–243.
- TIBSHIRANI, R. (1996): “Regression shrinkage and selection via the Lasso,” *Journal of the Royal Statistical Society, Series B (Methodological)*, 58(1), 267–288.

Appendix A Stochastic Search Variable Selection

A standard variant of the SSVS prior, similar to the one employed in [Piribauer \(2016\)](#) and [Piribauer and Crespo Cuaresma \(2016\)](#) for conventional spatial autoregressive models assumes each element of the parameter vector $\boldsymbol{\beta}$ to come from a mixture of two Gaussians,

$$\beta_r | \delta_r \sim \delta_r N(0, \underline{s}_{1r}) + (1 - \delta_r) N(0, \underline{s}_{0r}), \quad (\text{A.1})$$

for $r = 1, \dots, K$ with $\underline{s}_{0r} \ll \underline{s}_{1r}$ resulting in the spike-and-slab combination of components centered on zero. Variance hyperparameters are set for all β_r , that is $\underline{s}_j = (\underline{s}_{j1}, \dots, \underline{s}_{jK})'$ for $j \in \{0, 1\}$. We use the semi-automatic approach to scale the hyperparameters in \underline{s}_j by estimating the saturated model without applying shrinkage and obtain the posterior variance covariance matrix $\hat{\boldsymbol{\Sigma}}$. Subsequently, we set the spike component variances to $\underline{s}_0 = 0.01 \text{diag}(\hat{\boldsymbol{\Sigma}})$, and the slab component to $\underline{s}_1 = 100 \text{diag}(\hat{\boldsymbol{\Sigma}})$.

Moreover, we introduce $\boldsymbol{\delta} = (\delta_1, \dots, \delta_K)'$, which is a vector of latent quantities, binarily indicating which mixture component is active with respect to the r th coefficient β_r . This can be exploited to calculate posterior inclusion probabilities for all covariates, establishing a nexus to the model-averaging framework regarding interpretation. To estimate latent inclusion indicators, we hierarchically impose prior inclusion probabilities $\underline{\omega}_r$, chosen by the researcher, on the latent parameters

$$\boldsymbol{\delta} = \prod_{r=1}^K \omega_r^{\delta_r} (1 - \omega_r)^{1 - \delta_r}. \quad (\text{A.2})$$

[George and McCulloch \(1997\)](#) provide an intuitive interpretation for $p(\delta_r = 1) = 1 - p(\delta_r = 0) = \underline{\omega}_r$, which captures whether β_r is large enough to justify including the corresponding explanatory variable in the model. Arguably, a natural choice arising for ω_r ($r = 1, \dots, K$) is 0.5. The conditional posterior of $\boldsymbol{\delta}$ is a Bernoulli distribution, as presented by [George and McCulloch \(1993, 1997\)](#),

$$\begin{aligned} p(\delta_r = 1 | \bullet) &\sim \text{Ber}(\bar{u}_{1r} / (\bar{u}_{0r} + \bar{u}_{1r})), \\ \bar{u}_{1r} &= \underline{s}_{1r}^{-1} \exp(-\beta_r^2 / 2s_{1r}^2) \underline{\omega}_r, \\ \bar{u}_{0r} &= \underline{s}_{0r}^{-1} \exp(-\beta_r^2 / 2s_{0r}^2) (1 - \underline{\omega}_r). \end{aligned} \quad (\text{A.3})$$

The coefficient vector $\boldsymbol{\beta}$ is sampled using [Eq. \(8\)](#). Elements on the main diagonal of the prior variance-covariance matrix $\boldsymbol{\Sigma}$ whose latent indicator in $\boldsymbol{\delta}$ is equal to one are set to \underline{s}_{1r} , implying less prior influence and allowing for unrestricted variation. In the opposite case \underline{s}_{0r} is set, resulting in shrinkage towards zero.

Appendix B Bayesian MCMC estimation

Estimation is carried out running an Markov-chain Monte Carlo (MCMC) algorithm. One picks starting values for all parameters of the model and repeats the steps described below for a sufficient amount of iterations. The researcher decides on the number of iterations that depends on the mixing and convergence properties of the data at hand, and discards a number of initial draws as burn-in. Specifically, the sampler cycles through the following routine:

- (i) Given the values for parameters from the most recent iteration, draw σ^2 from the inverse Gamma distribution given in Eq. (12).
- (ii) This step depends on the specific shrinkage prior employed, items below are mutually exclusive:
 - **SSVS prior:** Draw δ from the Bernoulli distributed conditional posterior in Eq. (A.3), then update the prior variance-covariance matrix as described in the text and use it to sample from the Gaussian in Eq. (8).
 - **NG prior:** The local parameters ϕ_r are sampled independently from a generalized inverse Gaussian distribution⁹ as in Eq. (6), while the global shrinkage parameter λ comes from a Gamma distribution depicted in Eq. (7). Again, update the variance-covariance matrix $\text{diag}(\underline{\Sigma}) = \psi$ and sample the parameters using Eq. (8).
 - **DL prior:** Local and global scaling parameters are sampled from generalized inverse Gaussians given in Eq. (10) and Eq. (11), respectively. Draw the auxiliary variables T_r ($1, \dots, K$) mentioned above again from a generalized inverse Gaussian and set $\phi_r = T_r / \sum_{j=1}^K T_j$. As in the case of the SSVS and NG prior, the variance-covariance matrix must be updated $\text{diag}(\underline{\Sigma}) = (\varphi_1 \phi_1^2 \tau^2, \dots, \varphi_K \phi_K^2 \tau^2)'$ and used to sample from Eq. (8).
- (iii) Update ρ using a Metropolis-within-Gibbs step using the conditional posterior given in Section 3.

⁹We use the algorithm described and implemented by Hörmann and Leydold (2014) to sample from a generalized inverse Gaussian distribution.

Appendix C Data Appendix

Table C.1: Variables used in the empirical illustration

Variable	Description
Gross fixed capital formation	physical capital measured in terms of gross fixed capital formation (per capita). <i>Source:</i> Cambridge Econometrics
Initial income	Gross-value added divided by population, 2000. <i>Source:</i> Cambridge Econometrics
Physical capital	Gross fixed capital formation, 2000. <i>Source:</i> Cambridge Econometrics
Higher education workers	Share of population (aged 25 and over, 2000) with higher education (ISCED levels 1-2). <i>Source:</i> Eurostat
Lower education workers	Share of population (aged 25 and over, 2000) with lower education (ISCED levels 5-6). <i>Source:</i> Eurostat
Technology resources	Human resources in science and technology, share in persons employed, 2000. <i>Source:</i> Eurostat
Agricultural employment	Share of NACE A and B (agriculture) in total employment, 2000. <i>Source:</i> Cambridge Econometrics
Manufacturing & construction employment	Share of NACE C to F (mining, manufacturing, energy and construction) in total employment, 2000. <i>Source:</i> Cambridge Econometrics
Construction employment	Share of NACE F (construction) in total employment, 2000.
Market services employment	Share of NACE G to K (market services) in total employment, 2000. <i>Source:</i> Cambridge Econometrics
Output density	Gross-value added per square km, 2000. <i>Source:</i> Eurostat
Employment density	Employed persons per square km, 2000. <i>Source:</i> Eurostat
Population density	Population per square km, 2000. <i>Source:</i> Eurostat
Population growth	Average growth rate of the population for 1996-2000. <i>Source:</i> Eurostat
Labor force participation	Employed and unemployed persons as a share of total
Child dependency ratio	The ratio of the number of people aged 0-14 to the number of people aged 15-64, 2000. <i>Source:</i> Eurostat
Old-age dependency ratio	The ratio of the number of people aged 65 and over to the number of people aged 15-64, 2000. <i>Source:</i> Eurostat
Accessibility road	Potential accessibility road, ESPON space=100. <i>Source:</i> ESPON
Accessibility rail	Potential accessibility rail, ESPON space=100. <i>Source:</i> ESPON

Table C.2: European regions in the sample.

Country	Region
Austria	Burgenland, Salzburg, Kärnten, Steiermark, Niederösterreich, Tirol, Oberösterreich, Vorarlberg, Wien
Belgium	Prov. Antwerpen, Prov. Luxembourg, Prov. Brabant Wallon, Prov. Namur, Prov. Hainaut, Prov. Oost-Vlaanderen, Prov. Liège, Prov. Vlaams Brabant, Prov. Limburg, Prov. West-Vlaanderen, Région de Bruxelles-Capitale
Bulgaria	Severen tsentralen, Yugoiztochen, Severoiztochen, Yugozapaden, Severozapaden, Yuzhen tsentralen
Czech Republic	Jihovýchod, Severozápad, Jihozápad, Střední Čechy, Moravskoslezsko, Středné Morava, Praha, Severovýchod
Denmark	Hovedstaden, Sjælland, Midjylland, Syddanmark, Nordjylland
Estonia	Estonia
Finland	Åland, Länsi-Suomi, Etelä-Suomi, Pohjois-Suomi, Itä-Suomi
France	Alsace, Île de France, Aquitaine, Languedoc-Roussillon, Auvergne, Limousin, Basse-Normandie, Lorraine, Bourgogne, Midi-Pyrénées, Bretagne, Nord - Pas-de-Calais, Centre, Pays de la Loire, Champagne-Ardenne, Picardie, Corse, Poitou-Charentes, Franche-Comté, Provence-Alpes-Côte d'Azur, Haute-Normandie, Rhône-Alpes
Germany	Arnsberg, Leipzig, Berlin, Mecklenburg-Vorpommern, Brandenburg, Mittelfranken, Braunschweig, Münster, Bremen, Niederbayern, Chemnitz, Oberbayern, Darmstadt, Oberfranken, Detmold, Oberpfalz, Dresden, Rheinhessen-Pfalz, Düsseldorf, Saarland, Freiburg, Sachsen-Anhalt, Giessen, Schleswig-Holstein, Hamburg, Schwaben, Hannover, Stuttgart, Karlsruhe, Thüringen, Kassel, Trier, Koblenz, Tübingen, Köln, Unterfranken, Lüneburg, Weser-Ems
Greece	Anatoliki Makedonia, Thraki, Kriti, Attiki, Notio Aigaio, Dytiki Ellada, Peloponnisos, Dytiki Makedonia, Sterea Ellada, Ionia Nisia, Thessalia, Ipeiros, Voreio Aigaio, Kentriki Makedonia
Hungary	Dél-Alföld, Közép-Dunántúl, Dél-Dunántúl, Közép-Magyarország, Észak-Alföld, Nyugat-Dunántúl, Észak-Magyarország
Ireland	Border, Midlands and Western, Southern and Eastern
Italy	Abruzzo, Molise, Basilicata, Piemonte, Calabria, Bolzano-Bozen, Campania, Trento, Emilia-Romagna, Puglia, Friuli-Venezia Giulia, Sardegna, Lazio, Sicilia, Liguria, Toscana, Lombardia, Umbria, Marche, Valle d'Aosta, Veneto
Latvia	Latvia
Lithuania	Lithuania
Luxembourg	Luxembourg (Grand-Duché)
Netherlands	Drenthe, Noord-Brabant, Flevoland, Noord-Holland, Friesland, Overijssel, Gelderland, Utrecht, Groningen, Zeeland, Limburg, Zuid-Holland
Norway	Agder og Rogaland, Sor-Ostlandet, Hedmark og Oppland, Trondelag, Nord-Norge, Vestlandet, Oslo og Akershus
Poland	Dolnoslaskie, Podkarpackie, Kujawsko-Pomorskie, Podlaskie, Lodzkie, Pomorskie, Lubelskie, Slaskie, Lubuskie, Swietokrzyskie, Malopolskie, Warminsko-Mazurskie, Mazowieckie, Wielkopolskie, Opolskie, Zachodniopomorskie
Portugal	Alentejo, Lisboa, Algarve, Norte, Centro
Romania	Bucuresti-Ilfov, Sud-Muntenia, Centru, Sud-Est, Nord-Est, Sud-Vest Oltenia, Nord-Vest, Vest
Slovak Republic	Bratislavský kraj, Východné Slovensko, Stredné Slovensko, Západné Slovensko
Slovenia	Vzhodna Slovenija, Zahodna Slovenija
Spain	Andalucia, Extremadura, Aragón, Galicia, Cantabria, Illes Balears, Castilla y León, La Rioja, Castilla-la Mancha, Pais Vasco, Cataluña, Principado de Asturias, Comunidad de Madrid, Región de Murcia, Comunidad Foral de Navarra, Comunidad Valenciana
Sweden	Mellersta Norrland, Småland med öarna, Norra Mellansverige, Stockholm, Östra Mellansverige, Sydsverige, Övre Norrland, Västsverige
Switzerland	Central Switzerland, Northwestern Switzerland, Eastern Switzerland, Ticino, Espace Mittelland, Zurich, Lake Geneva
United Kingdom	Bedfordshire, Hertfordshire, Kent, Berkshire, Bucks and Oxfordshire, Lancashire, Cheshire, Leicestershire, Rutland and Northants, Cornwall and Isles of Scilly, Lincolnshire, Cumbria, Merseyside, Derbyshire and Nottinghamshire, North Yorkshire, Devon, Northern Ireland, Dorset and Somerset, Northumberland, Tyne and Wear, East Anglia, Outer London, East Riding and North Lincolnshire, Shropshire and Staffordshire, East Wales, South Western Scotland, Eastern Scotland, South Yorkshire, Essex, Surrey East and West Sussex, Gloucestershire Wiltshire, Tees Valley and Durham, North Somerset, Greater Manchester, West Midlands, Hampshire and Isle of Wight, West Wales and The Valleys, Herefordshire, Worcestershire and Warks, West Yorkshire, Inner London