Nonlinear impact estimation in spatial autoregressive models

Jean-Sauveur Ay, Kassoum Ayouba, and Julie Le Gallo

ULR CESAER, AgroSup Dijon, INRA, Université de Bourgogne Franche-Comté, 26 boulevard Petitjean, 21079 Dijon Cedex

August 12, 2018

Abstract

This paper extends the literature on the calculation and interpretation of impacts for spatial autoregressive models with an exogenous variable introduced in a nonlinear way. We show how the direct and indirect impacts can be computed, theoretically and empirically, in such a case. Rather than averaging the individual impacts, we suggest to plot them along with their confidence intervals. We also explicitly derive the form of the gap between impacts in the spatial autoregressive model and the corresponding model without a spatial lag and show that it is higher for spatially highly connected observations. We illustrate these results on the Boston dataset.

Keywords: Spatial econometrics, marginal impacts, spline, polynomials
JEL: C11, C21

1 Introduction

Spatial autoregressive (SAR) models are now widely used to analyze spatial economic interactions in various applications. For example, such a model is appropriate for the housing market since housing prices depend on prices of recently sold neighboring homes (Anselin and Lozano-Gracia, 2008). This dependence structure comes from the fact that sellers presumably use information on neighbooring homes to determine the asking price. The use of a SAR model, besides providing a richer characterization of the market, has important implications for impacts calculations (for early statements of this issue, see Anselin and Le Gallo (2006), Kim et al. (2003), Kelejian et al. (2006)). Indeed, their computation and interpretation is less straightforward than in standard multiple linear a-spatial regression models: any change in an explanatory variable for a given observation not only affects the observation itself (direct impact) but also affects all other observations indirectly (indirect impact).

LeSage and Pace (2009) show how to theoretically derive these marginal impacts in SAR models. For n spatial observations, they obtain a $n \times n$ matrix of impacts for one exogenous variable. In order to have a compact representation of these impacts, they propose to report one direct impact equal to the average of the diagonal elements of the matrix of marginal impacts and one indirect impact equals to the average row sums of the non-diagonal elements of that matrix. However, when the exogenous variable of interest is introduced in a nonlinear way in the SAR model (e.g. in the form of polynomial or splines function), averaging the impacts in such a way is irrelevant.

In this paper, we extend the work of LeSage and Pace (2009) on impacts computation and estimation when the exogenous variable of interest appears in a nonlinear way in a SAR model. We also derive the form of the gap between impacts in the spatial autoregressive model and the corresponding model without a spatial lag and show that it is higher for spatially highly connected observations.

The remainder of the paper is organized as follows. The second section presents the theoretical derivation of direct and indirect impacts associated with a nonlinear exogenous variable in a SAR model. The third section introduces the estimation strategy which is applied on the well-known Boston housing dataset described in the fourth Section. The obtained empirical results are presented in the fifth Section. The sixth section concludes and comments on possible extensions.

2 Impacts in theory

Consider the following spatial autoregressive model[1]

$$y = \rho W y + X\beta + f(z) + \epsilon \tag{1}$$

[1]Our approach can also be applied on spatial Durbin model.

where y is the $(n \times 1)$ dependent variable that exhibits variation across spatial observations, X is the $(n \times k)$ matrix of linear explanatory variables including an intercept term, with the associated parameters β contained in a $(k \times 1)$ vector, z is the $(n \times 1)$ variable the impact of which on y is nonlinear, and W is a specified constant $(n \times n)$ spatial weight matrix with the usual assumptions. This variable z of interest is of dimension 1 and is additively separable from the others X to simplify notations, the approch we propose can be easily extended to many variables introduced nonlinearly with interaction effects between them. We also assume that each term of the disturbance vector ϵ of dimension $(n \times 1)$ is normally and identically distributed with zero mean and variance σ^2. The scalar ρ measures the strength of the spatial dependence. $f(\cdot)$ is a linear-in-parameters function, for instance a polynomial function of degree p: $f(z) = \sum_{j=1}^{p} \gamma_j z^j$, a spline[2] function of order p and q knots: $f(z) = \sum_{j=1}^{p} \gamma_j z^j + \sum_{l=1}^{q} \delta_l (z - t_l)_+^p$.

Equation (1) can be rewritten in the following reduced form

$$y = V(W)X\beta + V(W)f(z) + V(W)\epsilon \tag{2}$$

with

$$V(W) = (I_n - \rho W)^{-1} = \begin{bmatrix} V_{11} & V_{12} & V_{13} & \cdots & V_{1n} \\ V_{21} & V_{22} & V_{23} & \cdots & V_{2n} \\ V_{31} & V_{32} & V_{33} & \cdots & V_{3n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ V_{n1} & V_{n2} & V_{n3} & \cdots & V_{nn} \end{bmatrix}$$

and I_n the identity matrix of order n. Note that we assumed that the matrix $I_n - \rho W$ is not singular to reach the reduced form in Equation (2).

In this paper, we are interested in the estimation of the partial derivative of y with respect to changes in the nonlinear variable z in our SAR model. In models containing a spatial lag, the measure of the partial derivative of the dependent variable with respect to an explanatory variable is less straightforward than in standard linear models. Indeed the standard linear regression interpretation of coefficient estimates ($\widehat{\beta}_q = \frac{\partial \hat{y}}{\partial x_q}$) as partial derivatives no longer holds in SAR model since the matrix of explanatory variable is transformed by the $n \times n$ inverse matrix $V(W)$. In such a model, any change to an explanatory variable for a given observation affects the dependent variable of the observation itself (direct impact) and potentially the dependent variable of all other observations (indirect impact) through $V(W)$. We elaborate on this observation to derive the impacts for the

[2]See Hastie and Tibshirani (1990) for details.

variable z introduced in a nonlinear way.

Starting from the reduced form in Equation (2), the matrix of responses to a change of the nonlinear variable z on y is given by

$$\begin{bmatrix} \frac{\partial y_1}{\partial z_1} & \frac{\partial y_1}{\partial z_2} & \cdots & \frac{\partial y_1}{\partial z_n} \\ \frac{\partial y_2}{\partial z_1} & \frac{\partial y_2}{\partial z_2} & \cdots & \frac{\partial y_2}{\partial z_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial y_n}{\partial z_1} & \frac{\partial y_n}{\partial z_2} & \cdots & \frac{\partial y_n}{\partial z_n} \end{bmatrix} = \begin{bmatrix} V_{11}f_z(z_1) & V_{12}f_z(z_2) & \cdots & V_{1n}f_z(z_n) \\ V_{21}f_z(z_1) & V_{22}f_z(z_2) & \cdots & V_{2n}f_z(z_n) \\ \vdots & \vdots & \ddots & \vdots \\ V_{n1}f_z(z_1) & V_{n2}f_z(z_2) & \cdots & V_{nn}f_z(z_n) \end{bmatrix}$$

with $f_z(z)$ the derivative of $f(z)$.

Since we have an $n \times n$ matrix of impacts, the challenge is to find a way to compactly report them. In the case of exogenous variables that enter linearly, LeSage and Pace (2009) suggest to compute the average of the main diagonal elements and the average of the off-diagonal elements of the impacts matrix to obtain a summary measure of the direct impact and the indirect impact respectively. This method of summarizing impacts, which should be seen as “the best we can do at the moment”, for a constant effect is irrelevant when the effect of a variable as z can be positive and negative at different parts of it support. Therefore, we propose to plot the individual impacts (i.e. impacts of each observation) along with their confidence intervals. The individual direct (*IDI*) and total (*ITI*) impacts of the nonlinear variable z on y_i are given by

$$IDI_i = \frac{\partial y_i}{\partial z_i} = V_{ii} f_z(z_i) \tag{3}$$

$$ITI_i = \frac{\partial y_i}{\partial z} = \sum_{j=1}^{n} V_{ij} f_z(z_j) \tag{4}$$

From these expressions, we can reach the individual indirect impact (*III*) as follows

$$III_i = ITI_i - IDI_i \tag{5}$$

3 Estimation strategy

To get the individual impacts for z, we proceed in two steps. We first estimate the parameters of the SAR model, and then in a second step use them to compute $V(W)$ and the derivative of $f(z)$, i.e. $f_z(z)$ for each observation i. Let us begin with the estimation of the SAR model. We follow LeSage and Pace (2009) who have shown how to estimate SAR models using a Bayesian methodology. We briefly recall the approach in the next paragraph.[3]

The Bayesian model or alternatively the posterior distribution of the model parameters up to a constant term is given by the product of the likelihood of the available data (y, X and z) and the prior distributions of the parameters of the model. To get the mean and variance for the model parameters, an option is to numerically integrate the posterior distribution. Unfortunately, the form of the posterior distribution is often complex, leading to intractable numerical integration. To circumvent this problem, we rely on the methodology known as Markov Chain Monte Carlo (MCMC) to estimate the SAR model parameters. The general idea behind MCMC is that rather than work with the posterior density of our parameters, the same goal could be achieved by analyzing a large random sample from the posterior distribution. Specifically, we use the *Metropolis within Gibbs sampling* since our posterior distribution can be split in three posterior distributions: one of unknown form for ρ and two for σ^2 and the parameters associated to the linear and nonlinear variables following a Gamma inverse and a Normal distributions respectively.

Once we have estimated the parameter of our model, to get the individual impacts, we need the derivative of $f(z)$, i.e. $f_z(z)$. A straightforward way of getting it is to analytically compute the derivative of $f(z)$. For instance, when $f(z)$ is a polynomial function of degree p, its derivative is given by: $f_{z_i}(z) = \sum_{j=1}^{p} j\gamma_j z^{j-1}$. However, in some cases (e.g. splines functions), using the analytical derivative of $f(z)$ may not be that simple and it is better to proceed differently. In this paper, we make use of simulations to get the derivatives. Starting from the SAR model, the derivative of $f(z)$ for observation i is given by

$$f_z(z_i) = \lim_{\xi \longrightarrow 0} \frac{\hat{y_{i2}} - \hat{y_{i1}}}{\xi} \tag{6}$$

where $\hat{y_{i1}}$ and $\hat{y_{i2}}$ are predictions of y_i from the model in Equation (1) using z and $z + \xi$ respectively. All the estimated parameters except those related to $f(z)$ are set to zero when calculating $\hat{y_{i1}}$ and $\hat{y_{i2}}$.

[3] A full presentation of the procedure can be found in LeSage and Pace (2009, pp. 124-153).

4 Data

We use the Boston dataset for the application. It is available to the R community as part of the library spdep (Bivand et al. (2013), Bivand and Piras (2015)). The dataset reports for each of the units of observation (506 in total), the (corrected) median value of owner-occupied homes in the Boston area ($CMEDV$), together with several variables which might help to explain its variation such as: per-capita crime rate ($CRIM$) by town, nitric oxides concentration (NOX), proportions of non-retail business acres per town ($INDUS$), proportions of units built prior to 1940 (AGE), weighted distances to five Boston employment centers (DIS), index of accessibility to highways (RAD), property-tax rate (TAX) per \$ 10,000, pupil-teacher ratios by town ($PTRATIO$); proportion of blacks (B) and percentage of the lower status of the population ($LSTAT$).

5 Results and discussion

We apply the methodology described above to Boston house prices. Specifically, we estimate the SAR specification in Equation (1), where y is the median value of owner-occupied homes ($CMEDV$), X the vector of the following linear variables: TAX, $INDUS$, $CRIM$ and AGE. The nonlinear variable z is represented by DIS. f is a cubic spline function used to capture the potential nonlinear effect of DIS on the dependent variable.

5.1 Results of the Bayesian estimation

We use the following starting values for the parameters of our Bayesian model: 0.5 for σ^2, 0.1 for ρ and 0.5 for each of the remaining parameters (MCMC starting values). We generate a total of 10,000 simulations, and then discard the first 1,000 as a "burn-in" – to mitigate startup effects. The Metropolis-Hastings algorithm used to draw values for ρ perform very well: the acceptance probability is 48%[4] The chains of retained simulations are presented in the appendix A and show no sign of nonstationarity. Note that we provide some more formal diagnostics to the visual (graphical) convergence check. Specifically, we monitor the convergence of our chains of parameters using diagnostics suggested by Geweke (1992) and Gelman and Rubin (1992). They provide no evidence against convergence for our chains.

[4] LeSage and Pace (2009) state that the optimal acceptance rate (i.e., the one which minimizes the autocorrelations across the sample values) is between 40% and 60%.

Table 1 presents point estimates (i.e. posterior means), associated measures of uncertainty (i.e. posterior standard deviations) and the upper and lower bound of a 95% highest posterior density interval[5] (columns labeled Q2.5% and Q97.5%) for the parameters of our model.

Table 1: Parameter estimates

	MEAN	STD	Q2.5	Q97.5
CONST	22.8258	2.3160	18.2907	27.4274
SPLINES1	-13.4145	3.4085	-20.1183	-6.7841
SPLINES2	-9.3020	3.2576	-15.7103	-2.8356
SPLINES3	-11.8411	4.1098	-19.7626	-3.9209
CRIM	-0.1101	0.0342	-0.1781	-0.0432
TAX	-0.0051	0.0022	-0.0094	-0.0007
AGE	-0.0595	0.0131	-0.0847	-0.0336
INDUS	-0.3057	0.0603	-0.4241	-0.1863
σ^2	25.9011	1.6983	22.7676	29.4221
ρ	0.6890	0.0290	0.6285	0.7448

We wish to compare the Bayesian MCMC sampling procedure and the maximum likelihood method results.[6] To this end, the maximum likelihood estimates for the SAR model and sample data are presented in Table 2 alongside Bayesian estimates based on MCMC sampling scheme. We also report the t-statistics for the two estimations methods as in LeSage and Pace (2009, pp. 143-145). We observe that all estimates and t-statistics are nearly identical, suggesting they would produce similar inferences. It appears that all coefficients and the spatial lag parameter are significant.

[5]This coverage region is analogous to the 95% confidence intervals.

[6]The maximum likelihood estimation results are available in the appendix B.

Table 2: Comparison of SAR model estimates

	Max. Likelihood		Bayesian M-H	
	Estimates	t-stat	Estimates	t-stat
CONST	22.5830	9.1977	22.8258	9.8556
SPLINES1	-13.3724	-3.9837	-13.4145	-3.9355
SPLINES2	-9.2372	-2.8515	-9.3020	-2.8554
SPLINES3	-11.7986	-2.9195	-11.8411	-2.8811
CRIM	-0.1086	-3.2452	-0.1101	-3.2144
TAX	-0.0050	-2.2547	-0.0051	-2.2828
AGE	-0.0590	-4.4745	-0.0595	-4.5202
INDUS	-0.3029	-5.0132	-0.3057	-5.0650
σ^2	25.2490		25.9011	
ρ	0.6949	22.805	0.6890	23.6795

5.2 Impacts estimation

For each of the 10,000 iterations, we calculate the individual direct, total and indirect impacts of the variable *DIS* introduced in a nonlinear way in the specification. Let us focus on the individual direct impacts to clearly get how they are calculated. For observation i, we compute for each iteration the individual direct impact for the corresponding value of *DIS* using the parameter estimates and the derivative of the nonlinear function f, with equation (6). We end up with a chain of individual directs impacts corresponding to the number of iterations. As above, we discard the first 1,000 draws. On the remaining chain, we calculate the mean and the 2.5% and 97.5% percentiles which correspond respectively to the estimate of the individual direct impacts and its 95% confidence interval. The procedure is repeated for each of the 506 observations of *DIS*. The estimation of the individual total and indirect impacts follows the same procedure.

Figure 1 plots for each value of *DIS*, the direct and indirect impacts along with their confidence intervals. In the top panels of Figure 1, we plot the direct and indirect impacts of *DIS* for each observation i. We have also applied a LOESS regression to these two sets of data point to get some more readable curves of impacts (bottom panels)

Figure 1: Representation of direct and indirect individual impacts

Two important observations follow from Figure 1. First, we observe that the impacts (direct and indirect) are negative, suggesting that Boston employment centers are attractive: the closer a house is to city center, the higher its price and the price of its neighbors. Second, the marginal impacts are nonlinear: the more *DIS* increases, the less are the magnitude of the impacts. This means that an increase of *DIS* of one mile has a higher impact on values of homes close to the center compared to the others. The non-linearity of the impacts implies that taking their average is indeed irrelevant in our case, and may lead to erroneous conclusions.

To better see the nonlinear effect of *DIS*, we directly plot it against the dependent variable *CMED*.[7] For that purpose, for the value of *DIS* for observation i, we calculate for each iteration its prediction using the parameter estimates. Note that we use the unbiased predictor proposed by Kelejian and Prucha (2007, page 367) which is compatible with our SAR model in Equation (2): the prediction is based on the information set $\Lambda = \{Z, W\}$ where Z is the vector of linear and nonlinear exogenous variables. To get rid of the effect of the other exogenous variable and focus only on *DIS*, we set as in Fox (1987) their values to their mean. From there, we follow the same procedure as in the calculation of marginal impacts to end with an estimate of the effect of distance and the 95% confidence interval. The procedure is repeated for each of the 506 values of *DIS*. The obtained plot (see Figure 2) clearly shows the nonlinear effect of *DIS* on the endogenous variable *CMED*.

Figure 2: Impact of *DIS*

5.3 Structure of W and impacts in nonlinear models

In this section, we analyze the link between the spatial structure of the observations and the form of the impacts in nonlinear models in the context of the application. To that purpose, let us begin with the individual direct impacts. From Equation (3), we know that when there is no spatial autoregressive term (i.e. when the model is a-spatial), the individual direct impact is $f_z(z_i)$. Therefore, the difference between a-spatial and spatial models individual direct impacts is the following

$$GAP_i^{Direct} = f_z(z_i)\,[1 - V_{ii}] \tag{7}$$

Two observations emerge from the expression in Equation (7). First, the gap between the two individual direct impacts depends on the structure of spatial connections between observations ($w_{ij}, i, j = 1, ..., n$). Second, from the Leontief expansion, the elements of $V(W)$ are always positive since the estimate of the spatial autoregressive coefficient is positive. Therefore, the sign of the gap between

[7]The coefficients associated to splines variables in Table 1 are not directly interpretable.

the individual impacts in the a-spatial and the spatial model is of the sign of $f_z(z_i)$. As a consequence, $f_z(z_i)$ is the upper bound (when $f_z(z_i)$ is negative) or the lower bound (when $f_z(z_i)$ is positive) estimation of the individual impacts in the spatial autoregressive model.

Let us move now on individual total impacts. From Equation (4), we can deduce that for an a-spatial model the individual total impact is $f_z(z_i)$, which is the same as the individual direct impact.[8] Therefore, the difference between a-spatial and spatial individual total impacts is the following

$$GAP_i^{Total} = f_z(z_i)\,[1 - V_{ii}] + \sum_{j=1, j\neq i}^{n} V_{ij} f_z(z_j) \tag{8}$$

The aforementioned observations, made for individual direct impacts, are also valid for individual total effects. In addition, one can observe that the magnitude of the gap between the individual total impacts is higher than that of individual direct impacts: $\mid GAP_i^{Total} \mid > \mid GAP_i^{Direct} \mid$.

To confirm these observations, we plot for each observation i the direct and total impacts against $f_z(z_i)$, which is the total impact for the a-spatial model. The procedure is repeated for each observation in the Boston dataset and the obtained results are plotted in Figure 3.

Figure 3: Impact in the a-spatial model and in the SAR model

The analysis of Figure 3 confirms our previous observations. First, the derivative $f_z(z_i)$, indeed appears as the upper bound estimation of the total impact. This means in our case that the estimated total impacts in the a-spatial model are overestimated if the data generating process is a SAR model. Second, the gap is higher in absolute value for total impacts compared to direct ones. Third, the gap between the two curves is increasing as we move towards the center (low distances). Since the observations close to the center are often the most connected (dense zones), that suggests that the more connected is the observation, the more is the gap (i.e. the gap depends on spatial connections among observations).

To confirm this intuition, we focus on the structure of the spatial weights matrix. Specifically, we calculate for each observation, its number of neighbors defined as the number of houses in a 1,000 meters buffer. Then we plot for each observation the gap between $f_z(z_i)$ and the total impact against its number of neighbors. The result is presented in Figure 5 and clearly shows that central observations are

[8]This makes sense as there is no individual indirect impact in a-spatial models.

associated with higher gap. This result is interesting as it states that the more an observation is linked to the others, the more important is the bias if spatial autocorrelation is wrongly omitted. In the Boston housing market, our result states that median housing prices decrease more rapidly for houses located in dense areas than for the houses located in sparse areas.

This result has another important implication: it confirms that averaging the impacts may overlook important heterogeneity even if the exogenous variable of interest is linear. Indeed, for a linear exogenous variable x_k with the associated parameter β_k, the individual direct impact is $V_{ii}\beta_k$. If β_k is for example positive, only considering the average leads to the underestimation of the impact of x_k for observations in central areas and the overestimation in sparse areas. Therefore, in specific applications where considering the heterogeneity of impacts is important, for instance to design place-based policies, it is better to take advantage of the "noise" introduced by spatial autocorrelation between observations rather than smoothing the impacts by averaging them, even for linear exogenous variables.

Figure 4: Relation between the number of neighbours and the GAP

6 Conclusion

We provide in this paper a framework for the estimation of impacts associated with an exogenous variable introduced in a non linear way in spatial autoregressive model. We also show that instead of averaging the impacts, one should consider to plot them along with their confidence interval. Indeed, averaging the impacts, besides being inappropriate for nonlinear exogenous variable, smooths spatial interaction effects which may be of interest in spatial autoregressive models.

References

Anselin, L. and Le Gallo, J. (2006). Interpolation of air quality measures in hedonic house price models: spatial aspects. *Spatial Economic Analysis*, 1(1):31–52.

Anselin, L. and Lozano-Gracia, N. (2008). Errors in variables and spatial effects in hedonic house price models of ambient air quality. *Empirical economics*, 34(1):5–34.

Bivand, R., Hauke, J., and Kossowski, T. (2013). Computing the jacobian in gaussian spatial autoregressive models: An illustrated comparison of available methods. *Geographical Analysis*, 45(2):150–179.

Bivand, R. and Piras, G. (2015). Comparing implementations of estimation methods for spatial econometrics. *Journal of Statistical Software*, 63(18):1–36.

Fox, J. (1987). Effect displays for generalized linear models. *Sociological methodology*, pages 347–361.

Gelman, A. and Rubin, D. (1992). Inference From iterative simulation using multiple sequences. In *Bayesian Statistics*. J. Bernardo, J. Berger, A. Dawid and A. Smith, eds, Claredon Press, Oxford, pp. 457-511.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to calculating posterieur moments. In *Bayesian Statistics*. J. Bernardo, J. Berger, A. Dawid and A. Smith, eds, Claredon Press, Oxford, pp. 169-194.

Hastie, T. and Tibshirani, R. (1990). *Generalized additive models*. Wiley Online Library.

Kelejian, H. H. and Prucha, I. R. (2007). The relative efficiencies of various predictors in spatial econometric models containing spatial lags. *Regional Science and Urban Economics*, 37(3):363–374.

Kelejian, H. H., Tavlas, G. S., and Hondroyiannis, G. (2006). A spatial modelling approach to contagion among emerging economies. *Open economies review*, 17(4):423–441.

Kim, C. W., Phipps, T. T., and Anselin, L. (2003). Measuring the benefits of air quality improvement: a spatial hedonic approach. *Journal of environmental economics and management*, 45(1):24–39.

LeSage, J. and Pace, R. K. (2009). *Introduction to Spatial Econometrics*. CRC Press.

A Graphical convergence check

Figure 5: MCMC chains

B The maximum likelihood estimation results

Table 3: SAR estimation

	Dependent variable:
SPLINES1	-13.372^{***}
	(3.357)
SPLINES2	-9.237^{***}
	(3.239)
SPLINES3	-11.799^{***}
	(4.041)
CRIM	-0.109^{***}
	(0.033)
TAX	-0.005^{**}
	(0.002)
AGE	-0.059^{***}
	(0.013)
INDUS	-0.303^{***}
	(0.060)
Constant	22.583^{***}
	(2.455)
Observations	506
Log Likelihood	-1,574.987
ρ	0.69493
σ^2	25.249
Akaike Inf. Crit.	3,169.974
Wald Test	520.068^{***} (df = 1)
LR Test	303.378^{***} (df = 1)
Note:	*p<0.1; **p<0.05; ***p<0.01