Spatial Effects and Ecological Inference

Luc Anselin
Regional Economics Applications Laboratory
and Departments of Agricultural and Consumer Economics,
Economics, and Geography,
University of Illinois at Urbana–Champaign,
1301 Gregory Drive, Urbana, IL 61801
e-mail: anselin@uiuc.edu

Wendy K. Tam Cho
Departments of Political Science and Statistics,
University of Illinois at Urbana–Champaign,
361 Lincoln Hall, 702 South Wright Street, Urbana, IL 61801
e-mail: wendy@cho.pol.uiuc.edu

This paper examines the role of spatial effects in ecological inference. Both formally and through simulation experiments, we consider the problems associated with ecological inference and cross-level inference methods in the presence of increasing degrees of spatial autocorrelation. Past assessments of spatial autocorrelation in aggregate data analysis focused on unidimensional, one-directional processes that are not representative of the full complexities caused by spatial autocorrelation. Our analysis is more complete and representative of true forms of spatial autocorrelation and pays particular attention to the specification of spatial autocorrelation in models with random coefficient variation. Our assessment focuses on the effects of this specification on the bias and precision of parameter estimates.

1 Introduction

For at least half a century, it has been known that making accurate individual-level inferences from aggregate data is extremely difficult (Robinson 1950; Goodman 1953, 1959). The problem, commonly called the ecological inference problem, is complex and multifaceted, and our understanding of the nuances involved in modeling aggregate data is continually evolving. With any set of aggregate data—and with real-world data, in particular—it is

Authors' note: L. Anselin's research was supported in part by NSF Grant BCS–9978058 to the Center for Spatially Integrated Social Science (CSISS). Thanks go to Brian Gaines, Michael Ward, the participants in the QMSS Colloquium at the University of California, Santa Barbara, and three anonymous referees for many helpful suggestions, to Gary King for providing the spatial weights matrix from his simulation experiments, to Michele Casper and Joel Halversen for providing access to the stroke data, and to Simon Jackman and tamarama for lending computing power. All references to unpublished material are available at the *Political Analysis* Web site.

difficult to ascertain whether the conditions under which micro-level or individual-level data will aggregate consistently and without bias are met. Indeed, making such determinations often proves sufficiently difficult to render the desired data analysis infeasible.

To gain some insight into the ecological inference problem, we approach the estimation from an undertraveled path. In particular, we focus on inroads that can be made from an explicit consideration of spatial effects (i.e., spatial autocorrelation and spatial heterogeneity) and how they affect the information contained and observable in aggregate units. We demonstrate that the studies of spatial effects thus far in this context are inadequate, in both breadth and depth. Although it is extremely difficult, if not impossible, to characterize all forms of spatial effects, there is much room for progress. In this article, we simulate several realistic forms of spatial effects and then assess the impact of these effects on the ability to make reliable cross-level inferences.

We proceed by first defining spatial heterogeneity in the context of the ecological inference problem. Next, we discuss the interrelationship between spatial effects and aggregation bias. We follow this with a brief outline of King's (1997) "solution" to the ecological inference problem before proceeding to classify different types of spatial effects and to discuss how they pertain to the problem of making cross-level inferences. We examine spatial effects in an actual data set that conforms to the standard ecological inference design and then demonstrate the role of different degrees of spatial autocorrelation through a series of Monte Carlo simulations. In the simulations, we focus on the bias and precision of ecological inference estimators. We conclude by summarizing the impact of spatial effects on aggregate data analysis and outlining some directions for future research.

2 Extreme Spatial Heterogeneity

Thus far, much of the work on the ecological inference problem in the mainstream social science literature has focused on the condition of aggregation bias, while other definably troublesome characteristics of aggregate data have received less attention (e.g., Ansolabehere and Rivers 1997, King 1997; Cho 1998). The role of spatial effects, in particular, has not been examined extensively or has been dismissed as less consequential than the aggregation bias assumption. Geographers depart from this general mindset in their study of the "modifiable areal unit problem" (MAUP) (Openshaw and Taylor 1979), a problem that is isomorphic to the ecological inference problem. While geographers are also concerned with aggregation bias, they somewhat uniquely place considerable emphasis on the issue of "zoning." That is, they are mindful of the spatial arrangement of the data and the specific size of the aggregate units of observation. In this paper, we meld this geographic perspective with the literature on ecological inference by considering the ecological inference problem in the context of spatial effects. For example, ecological inference can be seen as an example of spatial heterogeneity, or the phenomenon whereby a model (i.e., parameters, functional specification, error specification, etc.) is not constant across spatial observations (Anselin 1988, 1990). This is distinct from spatial autocorrelation, which is the match between attribute similarity and locational similarity, i.e., when large or small values coincide in space (positive spatial autocorrelation).[1]

For ease of exposition, we use the same notation as King (1997) and couch our examples in the same terminology and setting, that of racial voting patterns in geographic units called "precincts." However, it should be clear that this setup does not limit the generality of our

[1]The classic reference on spatial autocorrelation is Cliff and Ord (1981). We refer the reader to this source for a more technical treatment of spatial autocorrelation and spatial autocorrelation statistics.

treatment, and the overall framework applies equally well to other settings where ecological regression may be performed, such as the one considered in our empirical example.

We begin with the basic accounting identity that relates the total rate of voter turnout by precinct, T_i, to the composition of the precinct's population with respect to two mutually exclusive and exhaustive subgroups, say, proportion nonwhite, X_i, and proportion white, $(1 - X_i)$. The parameters in the model are the unknown (and usually unobservable) proportions of nonwhite turnout, β_i^b, and of white turnout, β_i^w. The relationship between these variables in the accounting identity is

$$T_i = \beta_i^b X_i + \beta_i^w (1 - X_i). \tag{1}$$

This accounting identity holds exactly for each of the p precincts in the data set, yielding a system with p equations (one for each precinct) and $2p$ unknowns (two parameters for each precinct).

From a classical (non-Bayesian) perspective, estimation of the parameters in this model is a special case of the *incidental parameter problem*. More precisely, no consistent estimator can be constructed for the individual parameters, since no informational gain results from obtaining further observations. Instead, each new observation creates two additional parameters to estimate. The standard approach for dealing with this issue is to treat the incidental parameters as "nuisance parameters" and to condition the estimation process on their values to obtain consistent estimators for the other (nonnuisance) parameters of interest [for a review of the technical issues, see Lancaster (2000), and the classic paper by Neyman and Scott (1948)]. However, this approach does not apply in the ecological inference context, since the parameters of interest are, in fact, the incidental parameters. Thus, in a classical framework, there is no consistent estimator that can be constructed for the individual parameters in an ecological inference problem, β_i^b and β_i^w.

In the spatial econometric literature, this situation is referred to as *extreme spatial heterogeneity* (Anselin 1988, 2000), and the "solution" is to impose spatial or geographical structure on the nature of the variation of the individual coefficients across observations. This approach, however, is only a partial solution, in the sense that the parameters to be estimated cannot be incidental and, thus, must be constrained to vary either continuously as a function of a small set of "fixed" parameters or in a discrete fashion by being constant across (spatial) subsets of the observations. Likewise, King applies this logic of parameter similarity in his model: The rationale underlying the EI estimator is to assume that the same set of means, β^b and β^w, underlies all observations in the data set. Any heterogeneity in the parameter values is modeled as random variation around these constant means. The constant means (and the covariance matrix for the errors around the means) are, then, the parameters that are estimated. Applying this random coefficient paradigm allows one to construct a statistical model to estimate the "optimal" predictors for the individual parameters, β_i^b and β_i^w, based on the estimates of the overall mean and the associated covariance matrix (e.g., Griffiths 1972).

While it is possible that the imposition of a random coefficient structure will capture heterogeneity properly in some situations, the assumed structure is certainly not a panacea for all instances of aggregate data. Moreover, the random coefficient structure is neither general, flexible, nor robust, and furthermore, it ignores other more definite forms of spatial structure that may be the source of the heterogeneity (Cho 1998; Freedman et al. 1998; Anselin 2000). A critical issue here, and one that has been addressed sparsely, at best, is determining the extent to which the assumed spatial structure can be verified solely through observations of the aggregate data. We touch on this difficult issue but focus most of our attention on the effects that arise from different forms of spatial effects in aggregate data.

3 Spatial Effects and Aggregation Bias

Before turning to specific models, we review the connection between spatial effects and the more familiar notion of aggregation bias. It is important to note that while we artificially separate these two conditions in our simulated data, aggregation bias and spatial effects usually coexist in real data (King 1997, p. 159; Cho 1998). Hence, our analysis of the consequences of spatial effects are conservative estimates, since the quality of the estimates is likely to be reduced by additional and simultaneously appearing sources (i.e., aggregation bias) as well.

Recent work on the ecological inference problem has emphasized that the origin of the cross-level inference problem is aggregation bias. In this context, aggregation bias is said to occur when the parameters in the model are correlated with the regressors.[2] For instance, for the model outlined above, if the parameters, β^b and β^w, are correlated with the regressors, X, then aggregation bias exists, and making cross-level inferences is not straightforward. In terms of the random coefficient model that we consider in more detail below, aggregation bias amounts to a correlation between the random variation around the common mean and the regressor X. As long as the latter is assumed to be exogenous, there cannot be any aggregation bias. However, the assumption of exogeneity is suspect and not typically reasonable. Indeed, one can conceive of many instances where the variation in either β_i^b or β_i^w would, in fact, be a function of X. Hence, to produce accurate estimates of the micro-level behavior, one must respecify the model in such a way that the parameters will be mean independent of the regressors. One way this can be accomplished is by incorporating the functional relation between the parameters and the X explicitly.

Since problems of ecological inference typically involve geographic units, the link between spatial patterns and ecological inference should be evident. The nature of the effect, whether one exists, and how different forms of spatial effects affect the analysis, however, are not clear. Achen and Shively (1995, chap. 4) focus much of their discussion on a problem they term "intraconstituency spatial autocorrelation—the unmeasured similarity of voters in the same district." Their claim is that a properly specified aggregate data model must control for intraconstituency spatial autocorrelation. In making this claim, they may seem to deviate from the literature that places aggregation bias at the forefront of conditions that must be controlled. However, while the Achen and Shively claim seems somewhat unorthodox at first, a closer reading reveals that they consider spatial autocorrelation and aggregation bias to be virtually one and the same. For instance, they state that "[d]ifferent constituencies will exhibit different loyalty and defection rates, and it is only by quirk that these differences will fail to correlate with the aggregate independent variable" (1995, p. 106). The implication is clearly that spatial autocorrelation is symptomatic of aggregation bias. Achen and Shively's (1995, p. 114) discussion of a solution—"what is needed for most applications is strong substantive knowledge of how individuals group themselves into constituencies and how best one might control for the resulting differences in mean disturbances"—while framed in terms of spatial autocorrelation, clearly implies that fixing the problem of spatial autocorrelation will simultaneously aid in alleviating aggregation bias.

[2] In econometrics, aggregation bias is not necessarily identified solely with this correlation but can also pertain to aggregation across functional forms, etc. We focus here on the usual interpretation in the ecological inference literature and leave the more complex aggregation issues for future work.

A key observation, then, is that the problems of spatial autocorrelation and aggregation bias, while separate and fundamentally different conditions, are, paradoxically, interrelated. It is beyond the scope of the current paper to delve further into the link between the two obviously related assumptions in aggregate data analysis. However, given the clear connection, it remains an important avenue of future research. Here, we simply draw attention to the intricate interconnection in real data but now shift our focus to simulated data to explore the consequences of alternative conceptualizations of some forms of realistic spatial autocorrelation and spatial heterogeneity (in the absence of aggregation bias).

4 King's Ecological Inference Solution

The ecological inference solution proposed by King can be viewed as a combination of a random coefficient approach and the familiar method of bounds, couched primarily in a Bayesian framework. It is useful for expository purposes to characterize the EI estimator primarily as a random coefficient model and to set aside the role of the bounds for now.[3] A basic assumption is that the heterogeneity in the model is due to random variation around an underlying common mean, which yields a regression model with heteroskedastic error terms. Formally, $\beta_i^b = \beta^b + \varepsilon_i^b$ and $\beta_i^w = \beta^w + \varepsilon_i^w$, such that

$$T_i = \beta^b X_i + \beta^w (1 - X_i) + u_i, \tag{2}$$

with $E(u_i) = E[\varepsilon_i^b X_i + \varepsilon_i^w (1 - X_i)] = 0$, provided that $E(\varepsilon_i X_i) = 0$. The latter condition is satisfied as long as the X_i are exogenous, i.e., as long as there is no aggregation bias as defined above. The variance term then becomes $\mathrm{Var}(u_i) = \sigma_b^2 X_i^2 + \sigma_w^2 (1 - X_i)^2 + 2\sigma_{bw} X_i (1 - X_i)$, where the X_i are again exogenous, which is heteroskedastic as long as the X_i vary across precincts. The extent to which this heteroskedasticity matters in any given situation is largely an empirical matter—it depends on the heterogeneity among the X_i, which can be tested by means of standard regression diagnostics (e.g., using a Breusch–Pagan Lagrange multiplier statistic).

Note that in contrast to much of the theory and practice in the random coefficient literature, the error covariance in the EI estimator, σ_{bw}, is taken to be nonzero, since the accounting identity implies an exact linear relationship between the two parameters.[4] King exploits this relationship in a model diagnostic he calls a "tomography plot." This plot allows one to visualize the constraints on the acceptable parameter pairs provided the original ecological inference problem is exactly two-dimensional. Since the parameters are probabilities, the plot is immediately constrained to the unit square. Each line in the plot embodies the logically possible parameter values. Given a value for β^b, the value of β^w is determined through the accounting identity. Hence, the possible parameter pairs, (β^b, β^w), lie on a line.

As long as there is a common underlying mean, or as long as the precinct-specific bounds do not logically preclude the existence of a common mean, this mean can be estimated consistently without any further distributional assumptions. Specifically, a consistent estimator

[3] Setting aside consideration of the role of the bounds does not have a material consequence on our discussion. The main role of the bounds is to yield a truncated normal density as the basis for the likelihood and to provide additional information for use in the derivation of the posterior density for each individual coefficient. Besides the inclusion of the method of bounds, the other features of the EI approach are standard to random coefficient estimation.

[4] The constraint is $\beta_i^b = \frac{T_i}{X_i} - \beta_i^w \frac{(1-X_i)}{X_i}$. Unless $X_i = 1, \forall i$, this implies a nonzero covariance between the two coefficients. In contrast, see, for example, Griffiths et al. (1979), where the covariance between the random components is set to zero.

such as feasible GLS does not require an assumption of normality. However, this does not hold when the truncation is explicitly considered as well. For example, the EI approach imposes normality to estimate the common means and covariances from a likelihood that does incorporate the parameter constraints in the form of a truncated bivariate normal.

Once the overall parameters are obtained, they can be used to construct optimal predictors for the individual coefficients. As shown by Griffiths (1972), in the standard random coefficient model, such an optimal predictor takes the form

$$\hat{\beta}_i = \hat{\beta} + \hat{\Sigma}\, x_i (x_i' \hat{\Sigma} x_i)^{-1} (y_i - x_i' \hat{\beta}), \tag{3}$$

where $\hat{\beta}$ is a vector of common means, $\hat{\Sigma}$ is a matrix that contains the estimates of the random error variances and covariances, x_i is a vector of observations on the explanatory variables, and $(y_i - x_i'\hat{\beta})$ is the residual for observation i. In other words, the best linear unbiased predictor is obtained by allocating the residual to each of the individual β_i, using weights that are a function of the value of the x_i and the covariance of the random errors. The result of employing such a specification is a model that yields a perfect fit for each observation.[5] Consequently, diagnostics based on the usual notion of fit or lack of fit are meaningless. The only notion of fit that may be used to construct diagnostics would be the one based on the overall mean, $\hat{\beta}$, but there is no observable counterpart to assess the properties of the individual $\hat{\beta}_i$. The EI estimator uses Bayesian constructs to derive the posterior distribution of the β_i, conditional upon the common parameters $\hat{\beta}$ and $\hat{\Sigma}$ [following the principles outlined by Griffiths et al. (1979)], while incorporating information on the precinct-specific bounds through T_i (using a truncated bivariate normal as the underlying distribution in the likelihood function).

A well-known practical problem in the estimation of random coefficient models is the lack of a positive definite covariance matrix, $\hat{\Sigma}$. There are a number of ad hoc solutions that have been proposed to deal with this problem. King avoids these issues altogether by enforcing positive definiteness (as well as an adherence to the parameter bounds) in the constrained maximum likelihood. Furthermore, in contrast to the classical treatment of prediction of the individual coefficients, King generates a full posterior distribution that incorporates, in addition to the common parameters, $\hat{\beta}$ and $\hat{\Sigma}$, the observation-specific bounds that are evident from the tomography plot as well.

To summarize, the EI solution attempts to account for a particular aspect of heterogeneity, conceptualized as random variation around a common mean. While this is a general method for dealing with heterogeneity, the method does not incorporate any information about the particular structure of the heterogeneity and is therefore not necessarily appropriate when employed on data that do not exhibit this specific form of heterogeneity. For instance, in some cases, instead of random variation around a common mean, the heterogeneity may be "spatial," i.e., distinct sets of geographic units will be heterogeneous from other distinct sets in the data (Cho 2001). The basic EI model makes no attempt to exploit or discover the underlying nature of this heterogeneity.[6] In this sense, the EI model is not a spatial model.

[5]Note that $y_i - x_i'\hat{\beta}_i = y_i - x_i'(\hat{\beta} + \hat{\Sigma}x_i(x_i'\hat{\Sigma}x_i)^{-1}(y_i - x_i'\hat{\beta})) = 0$.

[6]Note that our discussion here concerns primarily the basic EI model, i.e., the EI model with no covariates. We are well aware that one may employ the extended EI model wherein one can include covariates into the model to control for aggregation bias and spatial autocorrelation. If the correct covariates are introduced into the model, the spatial heterogeneity can be modeled. However, the "correct" covariates are typically not known and King supplies no formal method for choosing them. The sole passage in the book that is in the spirit of a "covariate test" recommends "walking around some of these neighborhoods, or standing by polling places, or reading the local press, or going to the supermarkets in the area" (King 1997, p. 281). King does not present a formal method for modeling heterogeneity.

5 Conceptualizing Spatial Models in the Aggregate Context

There are many ways to devise spatial models that can be categorized as different conceptualizations of a substantive ecological inference problem. We review some of these models below. We view these conceptualizations as a broad survey of the categories that embody the different types of heterogeneity that are evident in aggregate data. We focus on three basic forms of spatial/geographic aspects in data. The first conceptualization is based on spatially random coefficients. This category has the same general underlying rationale as the EI estimator, with the exception that spatial autocorrelation is introduced explicitly. The second specification that we review concerns spatial heterogeneity in the form of discrete changes in the underlying parameters. We call these spatial regime models. The last form involves continuous variation over space (spatially varying coefficients).

5.1 *Spatially Random Coefficients*

Incorporating full two-dimensional and multidirectional spatial dependence in a random coefficient specification such as EI, where different bounds constrain the parameter values of each observation, is not a straightforward process. Although King (1997) claims to assess the consequences of spatial effects, his Monte Carlo evidence is quite limited and characteristic of a time-series perspective toward spatial analysis. More specifically, his rendition of spatial autocorrelation is rooted in a unidimensional, one-directional process and, thus, is not generalizable to all, or even most, forms of "true" two-dimensional and multidirectional spatial autocorrelation. It is simple to see that in the spatial realm, precincts have multiple neighbors (multidirectional) and the relationship between two neighboring precincts is not one-sided (two-dimensional). Because his assessment of spatial effects is simplified, he does not account for some of the standard features commonly associated with spatial autoregressive processes such as simultaneity and induced heteroskedasticity (Anselin 1988).

Even our broadening of the assessment of the role of spatial effects in this paper is somewhat limited because of the multiplicity of ways in which they may be modeled. In addition, how this is implemented is highly consequential. Our treatment, in contrast to previous studies, touches the core of spatial effects, is the most comprehensive treatment to date, and is reasonably representative of several important forms of spatial effects.

Before proceeding to describe our assessment of the role of spatial effects, we briefly review King's simplified treatment of the same matter. In particular, King's specification, described in the context of Monte Carlo simulation experiments, is as follows:

$$\beta_i = \delta\beta_{i-1} + (1-\delta)u_i, \tag{4}$$

where u_i is a (draw from a) truncated normal variate and $\beta_1 = u_1$ (King 1997, p. 167). In this expression, King refers to δ as the "degree of autocorrelation." However, upon closer examination, it is obvious that this specification does not describe an autoregressive process but, rather, is a weighted average of two truncated random variables, which is more akin to a moving average specification. Moreover, this specification does not describe a process that is spatial and, thus, does not sufficiently serve as a model for assessing the effect of spatial autocorrelation. This alternative, nonspatial specification is not without consequences, since true spatial dependence precludes a recursive formulation for the process and instead requires all variates to be determined simultaneously. In other words, rather than taking β_1 as the initial starting point and then defining the other β_i's recursively, in a true spatial process, all the β_i are determined jointly and simultaneously.

In a random coefficient model, the spatial dependence pertains to the deviations around the common mean, or the observation-specific error terms. A spatial autoregressive process for these error terms implies that large or small deviations from the common mean will tend to occur in spatial clusters, rather than in a spatially random manner, as is assumed in the standard model. Formally,

$$(\beta_i - \beta) = \rho \sum_{j \neq i} w_{ij}(\beta_j - \beta) + \xi_i, \tag{5}$$

or

$$\varepsilon_i = \rho \sum_{j \neq i} w_{ij}\varepsilon_j + \xi_i, \tag{6}$$

where $\varepsilon_i = \beta_i - \beta$ are the error terms, ξ_i are i.i.d. innovation terms, and the index, j, refers to the "neighbors" of i, as defined by the nonzero elements, w_{ij}, of a spatial weights matrix W.[7] In matrix notation, with ε as the $p \times 1$ vector of random deviations from the common mean, β, W as the $p \times p$ matrix of spatial weights, and ξ as a $p \times 1$ vector of i.i.d. innovations, the usual spatial autoregressive process is given as

$$\varepsilon = \rho W \varepsilon + \xi \tag{7}$$

or as

$$\varepsilon = (I - \rho W)^{-1}\xi. \tag{8}$$

The inverse matrix represents the so-called spatial multiplier, which demonstrates the joint (simultaneous) nature of the spatial dependence. It is clear in this formulation that each ε_i is a function of all the ξ_j in the system through the elements of the matching row in the inverse matrix $(I - \rho W)^{-1}$. For ease of notation, we represent these elements by a_{ij}. The complex nature of the resulting covariance matrix for the random coefficient model can be seen if we substitute $\sum_j a_{ij}^b \xi_j^b$ for ε_i^b and $\sum_j a_{ij}^w \xi_j^w$ for ε_i^w, with a_{ij} as the row elements in the spatial multiplier inverse. This yields an error term of the form

$$u_i = \sum_j a_{ij}^b \xi_j^b X_i + \sum_j a_{ij}^w \xi_j^w (1 - X_i). \tag{9}$$

Since the ξ_j^b and ξ_j^w terms are uncorrelated across observations, the variance terms are still a function of the variance and covariance of ξ^b and ξ^w. However, the spatial multiplier terms, a_{ij}, induce extra heteroskedasticity. In addition, the autocorrelation yields nonzero covariance terms between the errors for different precincts. Ignoring this extra variance and covariance will yield inefficient estimates for the parameter, β, and biased estimates for the variance term.

It is important to note that the spatial autoregressive process of the "spatial lag" variety cannot be implemented in the EI model without violating the fundamental accounting

[7] By convention, the diagonals of a spatial weights matrix, or w_{ii}, are set to zero and the elements in each row are standardized so as to sum to one. For a recent and more complete review of the issues involved with specifying spatial weights and incorporating them into regression models, see Anselin and Bera (1998).

identity. For example, assume such a process in the dependent variable, T. For ease of notation, represent T as a vector, where

$$T = \rho WT + \beta^b * X + \beta^w * X^c. \tag{10}$$

Here the X and X^c are suitable vectors, β^b and β^w are vectors of precinct-specific coefficients, and the operation, $*$, is the direct (element by element) product. The corresponding reduced form is

$$T = (I - \rho\, W)^{-1}(\beta^b * X + \beta^w * X^c). \tag{11}$$

This is incompatible with the fundamental accounting identity because turnout in a precinct would then be a function of the racial turnout characteristics in all other precincts in the system, not just its own. In part, this is due to the assumption of an exogenous X_i. If one relaxes this assumption, one could conceive of the actual (spatially correlated) pattern of X_i to be the result of a spatial process yielding $(I - \rho W)^{-1} X^\ell$, where X^ℓ is the latent spatially random original layout.[8] However, since this clearly falls outside the standard approach of the accounting identity, we do not pursue it further. In our Monte Carlo simulations, we implement spatial autoregressive processes for the error terms.

5.2 *Spatial Regimes*

Spatial regimes consist of geographic subsets in which the model parameters assume distinct values (Anselin 1988). For example, one can think of two subregions of precincts, S_g and S_h, that together exhaust the district. Rather than allowing the parameters, β_i^b and β_i^w, to be different for each precinct i, they may take on two distinct values, say β_g^b for $i \in S_g$ and β_h^b for $i \in S_h$ (Cho 2001). The constancy of parameters across subregions is a testable hypothesis, for example, by means of a spatial Chow test (Anselin 1990). A crucial assumption in the spatial regime approach is the delineation of subregions. This should be exogenous to the model, or aggregation bias will occur. Ideally, the delineation and regime estimation should be carried out jointly (e.g., regimes with endogenous switching). However, this has not yet been attempted in a setting that also incorporates spatial dependence. In practice, exploratory spatial data analysis techniques such as local indicators of spatial autocorrelation (LISA) (Anselin 1995) may be useful in suggesting local clusters or spatial outliers that may form the core of a subregion. Also, the change point and parameter constancy literatures serve as auspicious starting points (Cho 2001).[9]

5.3 *Spatially Varying Coefficients*

Incorporating spatially varying coefficients is a hierarchical approach toward modeling the spatial variation of the model parameters across observations. This approach is also referred to as spatial expansion (Casetti 1997). In its simplest form, each precinct-specific

[8]This process should not be confused with a true spatial autoregressive (lag) model at the micro scale, i.e., a process that pertains to the individuals in each precinct. For example, such a process would be relevant if a voter made a decision on whether or not to vote dependent on the neighbors' decisions. This would represent a behavioral process that results in spatial autocorrelation at the subprecinct scale but is not observable at the precinct level itself.

[9]It should be noted that spatial regimes can (and often do) coexist with spatial autoregressive error processes. However, we do not pursue this avenue here.

parameter is a function of a (small) set of exogenous variables such as the terms in a spatial trend surface (a polynomial in the latitude and longitude of the locations). For example, the individual precinct estimates, β_i^b, may take the form $\beta_i^b = \beta^b + \gamma^1 z_i^1 + \gamma^2 z_i^2 + \varepsilon_i^b$. The presence of the random error leads to the same heteroskedastic disturbance that was evident in the random coefficient model, but the regression includes several additional terms as cross-products of the z_i^k and the X_i (for technical details, see Anselin 1992). As in the spatial regimes model, the constancy of the parameters is a testable assumption, and, in principle, the determination of the variables to be included in the expansion should be the subject of a careful specification search. In the context of ecological inference, additional complications arise because the parameters are bounded. Because of the bounds, the β_i^b that follows from the expansion specification must be contained in the 0–1 interval.

A recent variant of a spatially varying coefficient model is the geographically weighted regression (e.g., Fotheringham et al. 1998). Essentially, this is a form of spatial kernel estimation that may be useful as an exploratory technique or as a diagnostic to assess the presence of spatial heterogeneity. However, because it is not a model of that heterogeneity, it is not pursued further in the current context.

Here we focus simply on spatially random coefficients, leaving other forms of spatial models for future research. We mention the other forms above to highlight the vast array of specifications that would fall under the spatial rubric. This limited approach has obvious consequences for the generalizability of our results to the universe of possible problems from spatial effects in aggregate data. Nonetheless, it is more far-reaching than King's unidimensional, one-directional rendition of spatial effects and greatly expands the evidence on the performance of ecological inference estimators in the context of spatially autocorrelated data.

6 Stroke Mortality Rates at the County Level in Texas

Before embarking on an analysis via simulation, we present a real-world example of spatial effects. Our example comes from 1990 county-level data in Texas on incidences of deaths from stroke.[10] We consider both the data at the county level and the statewide aggregate. If we were trying to surmise what happens at the county level from the state-level data, this type of setting would be typical of the genre that we encounter in ecological inference problems. The data include, for each of 254 counties, the stroke mortality rate among whites (S), the proportion of white males (M), the proportion of white females (F), the total number of whites (W), the proportion of white males who died from strokes (β_m), and the proportion of white females who died from strokes (β_f). The relationship between these variables is

$$S = \beta_m M + \beta_f F. \tag{12}$$

Since the data set includes the proportion of white males who died from strokes (β_m) and the proportion of white females who died from strokes (β_f), we have the "truth" that one might seek through EI. In general, these variables will not be available. Their availability here is helpful because it allows us to assess the performance of various ecological inference estimators on this aggregate data set.[11]

This data set is particularly interesting for our study of spatial effects in aggregate data analysis because these data exhibit a high degree of spatial autocorrelation on a number of

[10] Data may be obtained from the authors upon request.

[11] Although these are not data on voting or some other political phenomenon, the relationship between this example and a more political example should be transparent.

dimensions. The stroke mortality rate among whites is highly spatially correlated, as are the proportions of males and females who died from strokes. The spatial autocorrelation of stroke mortality rates (S) is observable in aggregate data sets, but the spatial autocorrelation of the parameters is usually not observable. However, the spatial autocorrelation of the dependent variable is not as interesting for our quest here, since it does not necessarily imply spatial autocorrelation of the rates among males (β_m) and females (β_f). Instead, the spatial autocorrelation in S, the stroke mortality rate among whites, may result simply from spatial autocorrelation in M, the proportion of white males, a situation in which we are not primarily interested here (although the M are significantly spatially autocorrelated). Our main concern revolves around the spatial autocorrelation of the parameters, β_m and β_f.

Before discussing the spatial effects, we digress briefly to an exposition of measures of spatial effects. In our analysis, we employ a few standard measures of spatial autocorrelation. For instance, we examine Moran's I statistic to make assessments about the degree of global spatial autocorrelation in the data. Specifically, Moran's I statistic is

$$I = \frac{N}{S_0}\left(\frac{e'We}{e'e}\right), \tag{13}$$

where e is a vector of OLS residuals, W is the weights matrix, and $S_0 = \sum_i \sum_j w_{ij}$. This statistic can be thought of as a counterpart to the familiar Durbin–Watson statistic used to detect autocorrelation in time-series data. Spatial autocorrelation occurs when the similarity of values of interest is related to the locations of the units, i.e.,

$$\mathrm{Cov}(y_i, y_j) = E(y_i y_j) - E(y_i)E(y_j) \neq 0, \qquad \forall\, i \neq j. \tag{14}$$

It is clear from the specification of Moran's I that calculation of the statistic requires that we specify a weights matrix. In our analysis, we employ a number of weights. In our example here, our results incorporate both rook weights (four neighbors on average, common borders) and queen weights (eight neighbors on average, common borders and common vertices).

In addition to examining the global spatial autocorrelation statistic for our data, we also examine LISA statistics (Anselin 1995). This local Moran statistic is closely related to the global Moran's I; specifically, the average of the local I statistics is equal to the global I, to a factor of proportionality. Examining the local statistics allows us to identify observations that are "extreme contributions" to the global statistic by noting which values are, say, two or more standard deviations from the mean.

In our Texas stroke example, using rook weights, Moran's I statistic for β_m is 0.2519 and Moran's I statistic for β_f is 0.3966. Both are significant at the .001 level. With queen weights, Moran's I statistic is 0.2531 for β_m and 0.3818 for β_f. The significance level is likewise high when queen weights are employed. Figure 1 presents a visualization of local spatial autocorrelation. The plots on the left display the LISA statistics for the female stroke mortality rate, while the plots on the right display the LISA statistics for the male stroke rate.[12] The two plots on the top are based on rook weights, while the two plots on the bottom

[12]LISA statistics are "local indicators of spatial autocorrelation." Unlike Moran's I, which is a global indicator of spatial autocorrelation, LISA statistics provide an individual measure of local spatial autocorrelation for each observation in relationship to its defined neighbors. See Anselin (1995) for more details.

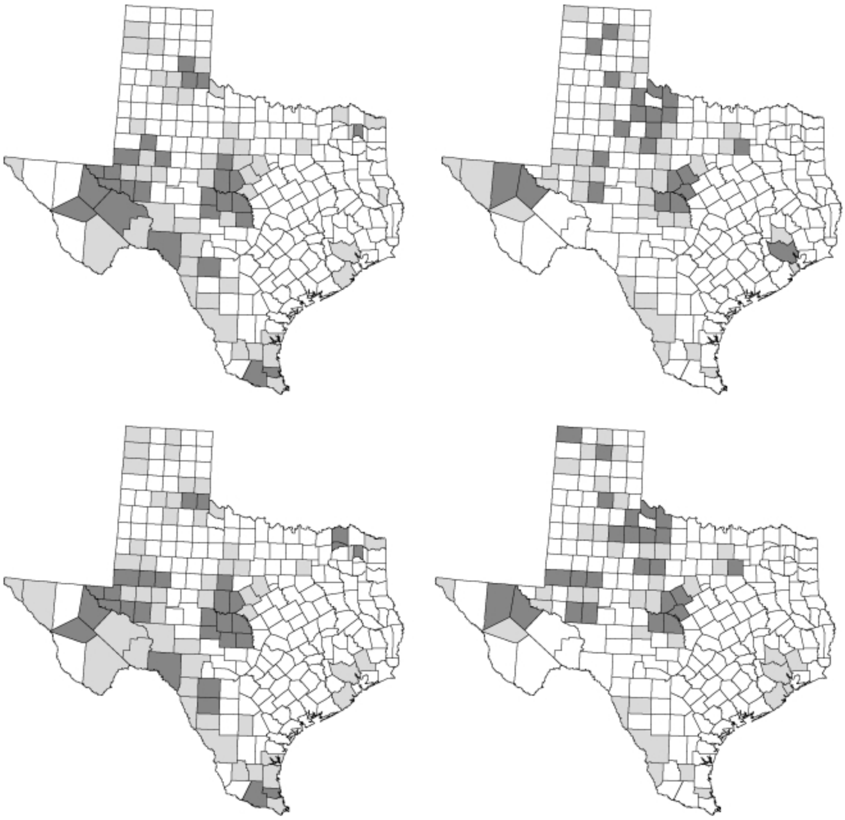

Fig. 1 LISA statistics. The LISA statistic is insignificant for the unshaded counties. The gray counties have LISA statistics significant at the .05 level. The darker-shaded counties have LISA statistics significant at the .01 level. The two plots on the top are based on rook weights. The two plots below are based on queen weights. The two plots on the left display LISA statistics for the female stroke mortality rate. The two plots on the right display LISA statistics for the male stroke mortality rate.

are based on queen weights. The unshaded counties have insignificant LISA statistics, while the shaded counties have significant LISA statistics. The lighter-shaded counties have LISA statistics that are significant at the .05 level. The darker-shaded counties have LISA statistics that are significant at the .01 level. In sum, both global and local statistics for spatial autocorrelation are highly significant in this data set, not only for the "observable" variables, S and M, but also for the unobservable β_f and β_m, the parameters of interest in an ecological analysis.

The results of estimating β_m, the male stroke mortality rate, and β_f, the female stroke mortality rate, via OLS and EI are reported in Table 1, where the truth is also reported. As we can see from comparing the model estimates to the true values, neither EI nor OLS performs very well on this data set.[13] Both models report estimates that are far from the

[13]The EI program comes with a number of diagnostics that one might employ in an analysis of these data. However, there are no diagnostics to examine spatial effects. Since King (1997) states in his book that spatial effects are inconsequential, one would not expect him to devote much effort toward developing diagnostics for these effects. Instead, most of the diagnostics are related to the aggregation bias assumption. From King's perspective, it is not useful to examine these diagnostics here, since he does not claim that aggregation bias causes or is an indicator

Table 1 Estimates of stroke mortality rates in Texas

	Males	*Females*
OLS/Goodman estimate	−0.065	0.081
	(0.009)	(0.008)
EI estimate	0.009	0.007
	(0.001)	(0.001)
Truth	0.0064	0.0095

Note. Standard errors in parentheses.

truth. Moreover, EI reverses the relationship between the coefficients. That is, although female stroke mortality rates are higher than male stroke mortality rates in Texas, the EI results conclude that male stroke mortality rates are higher than female stroke rates. Such a reversal of coefficient magnitude is noteworthy and would be instrumental in defining the resulting analysis. For instance, in Voting Rights cases, the parallel is that the EI model might report that whites voted in higher percentages for a certain candidate than black voters when the truth is exactly the opposite. The ruling by the judge would be highly affected (and wrongly so!) by such an assessment. So it is important to note EI's pivotal failure here. The OLS estimates may also be far from the truth and even out of bounds, but the relationship between the coefficients is correct. Moreover, the out-of-bounds estimates in this case is helpful because it immediately alerts the analyst to a problem. The EI estimate, on the other hand, while in the [0, 1] bounds, is incorrect on other dimensions (e.g., the relative direction of the two coefficients) and may leave the analyst with the faulty impression that the underlying truth has somehow been tapped. So the property of falling in-bounds in this case is somewhat of a liability, since it masks the erroneous estimate. In-bound estimates, then, which are typically considered to be an attractive characteristic of EI estimates, are not unambiguous benefits. We are, after all, not interested in *possible* estimates (i.e., estimates within the [0, 1] bounds) but in *correct* estimates that lead to *correct* inferences.[14]

While EI does not perform splendidly on this one example of a highly spatially correlated data set, it would be premature to draw definitive conclusions from what we observe in a single instance. These data exhibit spatial effects, but the spatial effects are not isolated in this example. As we have pointed out previously, it is not likely that spatial effects will be isolated in real data. Instead, they appear often with aggregation bias (King 1997; Cho 1998). In these data, aggregation bias is also present. The correlation between the parameter and the regressor is −0.43. While this is not an extremely high degree of correlation, the correlation is nonetheless present. Thus, we cannot surmise from this example what would happen if the spatial effects were isolated. Accordingly, to assess the effect of various forms of spatial effects on the performance of ecological inference estimators, we now turn to some Monte Carlo experiments.

---

of spatial effects. Instead, if the aggregation bias diagnostics gave one the impression that aggregation bias is not problematic in these data, King would advocate these estimates.

[14]One might be tempted here to compare a male stroke mortality rate of −0.065, which is impossible, with the EI estimate of 0.009, which is possible, but wrong. In this case, it is clear that aggregation bias exists (signaled simply from the out-of-bounds OLS estimate), and since EI is not robust to aggregation bias, one should clearly conclude that neither EI nor OLS is providing a reasonable estimate. King (1997, p. 282) himself cites out-of-bounds OLS estimates as clear evidence of aggregation bias and, further, acknowledges that basic EI is not robust to violations of the aggregation bias assumption. An in-bounds estimate is not synonymous with a correct or even reasonable estimate.

7 Some Monte Carlo Experiments

These experiments use the results of King's EI software both for the EI estimator itself and for Goodman's estimator (essentially OLS).[15] In this paper, we have limited our evaluation to OLS and the basic EI estimator and have not assessed any extensions of these estimators that explicitly account for the forms of spatial dependence and spatial heterogeneity described above.[16] This choice follows all of the previous research examining the effects of spatial dependence on the EI estimator.

Implementing spatial effects in simulation experiments that respect the fundamental accounting identity is not as simple as one might initially believe. Keeping values within the bounds requires some rule or procedure. In the few simulations in the literature (including King's), the precinct-specific values for the parameters β_i^b and β_i^w are drawn from a bivariate normal distribution with predetermined means and a given covariance structure. These draws are then converted to a truncated distribution by rejection sampling (i.e., the draws that do not fall within the 0–1 range are not used in the subsequent estimation). This procedure guarantees that the accounting identity is satisfied and that all values of T_i are within the zero–one bounds. However, when modeling spatial effects, care must be taken that the spatially transformed (autocorrelated) variates also still satisfy the accounting identity. Since spatial autoregressive processes tend to increase the variance of the underlying random variable, this condition must be checked explicitly. In other words, simply assessing that a set of random variates satisfies the truncation bounds does not guarantee that spatially transformed variates derived from them will fall within those bounds as well.

An additional complication results from explicit consideration of the spatial arrangement, or layout, of the data. We examine two layouts, one for a regular 10×10 grid system ($n = 100$) and the other for the arrangement of 179 countries referred to by King (2000) ($n = 179$).[17] These spatial layouts differ not only in the sample size but, more importantly, in the type of connectedness that is implied. On a regular grid, using the "rook" definition of neighbors (or, north, south, east, west), each cell has four neighbors, except for the boundary cells. The effect of those boundary cells varies with the sample size: for a 10×10 grid, the average number of neighbors is 3.6, with an average (row-standardized) weight of 0.28 (of the 100 grid cells, 64 have four neighbors, 32 have three, and four have two). In contrast, the connectivity structure of the King neighbors is very irregular.[18] While the average number of neighbors is only slightly higher than for the rook case (4.8), and the average weight is slightly less (0.21), this masks a high degree of underlying variability. One observation has as many as 16 neighbors, while 13 observations have only one neighbor, and 28 observations have only two neighbors. Controlling for all possible spatial layouts is

[15] All tuning parameters for the EI estimation software were kept at their default settings to avoid extra variability in the simulations and to keep some degree of comparability with King's results. We did, however, in the midst of our simulations, experiment with different settings. For example, the "Esim" global variable, the number of simulations used to generate the posterior distribution at the precinct level, is set to 100 as a default. We experimented with a value of 1000 for this parameter, but the results were essentially the same.

[16] See footnote 6 for an explanation of this choice.

[17] King's (1997) simulations used data sets of size $n = 100$ and $n = 1000$. He notes that his estimator performs better for larger data sets. However, our tests with two data sets ($n = 100$ and $n = 400$) did not indicate the same rise in performance. In any case, these results from his book are unimportant here, since King (2000) has agreed that the modeling of spatial dependence in his book is limited and not representative of true spatial effects, and his subsequent simulations do not indicate the sample size of his data sets or testing on data sets of different sizes.

[18] He seems to be using contiguity to define neighbors.

nearly impossible, so the two settings considered here should be considered to be illustrative only.[19]

Given the highly stylized nature of the design, we attempted to minimize the standard errors due to the simulation itself. Accordingly, we ran 1000 replications, which is four times more than for previously reported results in the literature. We also used the same base set of random variables in all of the spatial simulations (i.e., the spatial autoregressive parameter was applied to the same base set of random variates). In addition, to avoid variance instability resulting from the use of proportions/rates, each cell in the grid is constrained to have the same population base. Finally, to ensure that there is no aggregation bias in the simulated data, the X_i are constrained to be the same for all simulations and are generated as uniform random variates.[20]

We consider two designs, a truncated spatial design and a censored spatial design. We use these two approaches to ensure that the β_i are properly truncated.[21] The truncated model consists of rejection sampling applied to the full vector of errors, ε. Unlike the recursive unidimensional model considered by King, a "simultaneous" spatial process specifies the complete vector of variates for all precincts, so that a precinct-by-precinct rejection sampling is inapplicable. Only when a simulated vector of variates is obtained for which all of the ε_i are in the proper range can the result be retained. To implement this in a reasonable time, the standard error of the underlying untruncated normal distribution was lowered to 0.1.[22] The results using this approach are listed in Table 2. The censored model takes a slightly different perspective by allowing the error terms to follow a latent spatial autoregressive process that is *censored* at zero and one. In other words, the vector of ε and its matching β_i are generated as before, but all parameters that are less than zero are replaced by 0.0, and those greater than one are replaced by 1.0. While the resulting random coefficients can no longer be considered to be truncated bivariate normal, they are spatially correlated and satisfy the accounting identity. Arguably, this is a more realistic behavioral model for spatial dependence and fits into a rich tradition of latent variable models. In these simulations, we set $\sigma = 0.2$ to mimic the King setup (1997, p. 166). These results are listed in Table 3.[23]

In both cases, we follow King's design as closely as possible—the β_i^b and β_i^w are drawn from a truncated bivariate normal distribution by rejection sampling from an untruncated bivariate normal with parameters, $\beta^b = \beta^w = 0.5$, and correlation coefficient, $\rho = 0.3$ (see King 1997, p. 166). In the truncated case, the standard error, $\sigma_b = \sigma_w = 0.1$, while in the censored case $\sigma_b = \sigma_w = 0.2$. Note that King's specific choice of parameters here is particularly auspicious for subsequent estimation since both the untruncated and the truncated

[19]The King weights apparently were designed to mimic the arrangement of 179 countries in the world. While he did supply us with the weights matrix, he did not, however, supply any other details on its construction or the spatial layout of the data. Hence, we can provide no other details on his weights matrix.

[20]The latter feature may be unrealistic in the sense that empirical patterns for X_i show more regularity than a uniform random variate would suggest. On the other hand, it ensures the absence of aggregation bias and allows us to assess the specific consequences of spatial effects without the presence of other confounding circumstances.

[21]Note that it does not suffice to apply the spatial autoregressive transformation to truncated random variates such as ξ. The row sum in the spatial inverse matrix ranges from 1.25 for $\rho = 0.2$ to 5.0 for $\rho = 0.8$, which considerably increases the variance for the spatially autocorrelated variates. Without a further transformation, the latter tend to fall frequently outside the acceptable range.

[22]For example, with $\rho = 0.8$, the rejection sampling required the generation of one or two extra sets for roughly every third simulation. With the original standard error of 0.2, no acceptable vectors were generated for this value of ρ with fewer than 20 attempts, which rendered the procedure impractical. This highlights the interaction between the intrinsic variance of the underlying process and the resulting truncation.

[23]Note that in the censored case, it is not necessary to reduce the underlying variance to get a workable set of simulated values; this is done here only to keep comparability.

Table 2 Truncated model: Monte Carlo simulation results (1000 replications)

	Districtwide	*King's EI estimate*			*Goodman's regression (OLS)*		
ρ	β^b β^w	β_i^b β_i^w	*RMSE* *RMSE*	*MAE* *MAE*	β_i^b β_i^w	*RMSE* *RMSE*	*MAE* *MAE*
			Rook weights				
0.0	0.4998	0.4989	0.0155	0.0123	0.4990	0.0154	0.0122
	0.5000	0.5011	0.0164	0.0130	0.5011	0.0164	0.0129
0.2	0.4997	0.4989	0.0158	0.0126	0.4990	0.0158	0.0125
	0.5000	0.5010	0.0168	0.0133	0.5011	0.0167	0.0132
0.5	0.4996	0.4989	0.0178	0.0142	0.4990	0.0177	0.0141
	0.5000	0.5009	0.0189	0.0150	0.5010	0.0187	0.0148
0.8	0.4990	0.4984	0.0237	0.0190	0.4984	0.0234	0.0187
	0.4998	0.5007	0.0234	0.0200	0.5008	0.0248	0.0198
			King weights				
0.0	0.4999	0.5000	0.0110	0.0089	0.5000	0.0109	0.0088
	0.5002	0.5002	0.0116	0.0093	0.5002	0.0115	0.0092
0.2	0.4999	0.5000	0.0111	0.0089	0.5000	0.0110	0.0089
	0.5002	0.5001	0.0117	0.0093	0.5001	0.0116	0.0093
0.5	0.4998	0.5000	0.0119	0.0095	0.5001	0.0118	0.0095
	0.5002	0.5000	0.0126	0.0100	0.5000	0.0124	0.0100
0.8	0.4995	0.5000	0.0153	0.0121	0.5000	0.0150	0.0118
	0.5001	0.4996	0.0164	0.0130	0.4996	0.0161	0.0128

Note. In all cases, $\beta_i^b = \beta_i^w = 0.5$.

distribution will have the same mean. Moreover, due to the small variance, the degree of truncation will tend to be limited.[24] It is important to note as well that while our simulations are more realistic than King's simulations, the generation of data in this manner and to violate only the spatial assumption remain highly stylized. Hence, we should expect that the effects will necessarily be subtle.

In addition to the null case ($\rho = 0$), we considered three coefficient values for the spatial autoregressive process to simulate low, medium, and high spatial autocorrelation ($\rho = 0.2$, 0.5, and 0.8, respectively). The individual β_i^b and β_i^w coefficients are obtained through a three-stage process. First, draws are made from an untruncated bivariate normal distribution with mean zero and the covariance matrix specified above. Next, after the first step yields a vector of "error terms," ξ, these error terms are subsequently transformed into a vector of spatially autoregressive errors, ε, by setting $\varepsilon = (I - \rho W)^{-1}\xi$. Finally, we use the spatially autoregressive errors to compute the precinct-specific β_i coefficients by setting $\beta_i = \beta + \varepsilon_i$.

7.1 *Discussion of Monte Carlo Results*

Evaluating the performance of EI and OLS in the presence of spatial effects through our multiple sets of simulated data merits some special attention. In particular, unlike "traditional"

[24] King does not assess the possible effects from asymmetric truncation by allowing the means to differ, which is the more interesting case. In a Voting Rights context, for example, cases are brought before the court because of a strong prior that polarized voting exists, or that the true β^b is much different from the true β^w.

Table 3 Censored model: Monte Carlo simulation results (1000 replications)

ρ	*Districtwide* β^b / β^w	*King's EI estimate* β_i^b / β_i^w	*RMSE* / *RMSE*	*MAE* / *MAE*	*Goodman's regression (OLS)* β_i^b / β_i^w	*RMSE* / *RMSE*	*MAE* / *MAE*
			Rook weights				
0.0	0.5001	0.5016	0.0303	0.0241	0.5018	0.0305	0.0243
	0.5000	0.4993	0.0322	0.0257	0.4991	0.0323	0.0258
0.2	0.5009	0.5019	0.0313	0.0250	0.5019	0.0313	0.0249
	0.5001	0.4991	0.0330	0.0263	0.4991	0.0330	0.0264
0.5	0.5014	0.5020	0.0339	0.0271	0.5022	0.0349	0.0278
	0.5001	0.4995	0.0357	0.0286	0.4993	0.0366	0.0293
0.8	0.5031	0.5032	0.0351	0.0283	0.5036	0.0427	0.0340
	0.5003	0.5002	0.0373	0.0300	0.4997	0.0447	0.0358
			King weights				
0.0	0.4995	0.4996	0.0195	0.0156	0.4997	0.0201	0.0159
	0.4987	0.4987	0.0211	0.0168	0.4986	0.0217	0.0173
0.2	0.4995	0.4995	0.0198	0.0157	0.4997	0.0202	0.0161
	0.4984	0.4985	0.0214	0.0170	0.4983	0.0219	0.0175
0.5	0.4993	0.4996	0.0207	0.0164	0.4996	0.0216	0.0171
	0.4976	0.4975	0.0224	0.0179	0.4975	0.0234	0.0188
0.8	0.4993	0.4992	0.0222	0.0177	0.4998	0.0251	0.0199
	0.4951	0.4956	0.0240	0.0194	0.4950	0.0275	0.0223

Note. In all cases, $\beta_i^b = \beta_i^w = 0.5$.

simulation exercises, the proper standard of reference is not unambiguous. One difference that arises from evaluating the EI estimates in relation to the OLS estimates is that EI reports a districtwide estimate as well as precinct-specific point estimates (the mean of the posterior distribution) and the full distribution of the simulated posterior density. In contrast, the Goodman estimator provides the districtwide parameter with no variation in the precinct-specific parameters.[25] This difference is worth highlighting because one can compare either the overall estimates or, as we do, EI's precinct-level estimates and OLS's (single) districtwide estimate.

In our evaluation, we considered a number of criteria. First, to provide a sense for the effect of truncation and the sampling error of the simulations, we report the districtwide coefficients computed from the simulated "truths" for the individual precinct parameters, β_i^b and β_i^w. In most of our designs, the truncated and untruncated means should be very close, since there is no truncation effect on the mean with $\beta^b = \beta^w = 0.5$ for a symmetric distribution such as the normal distribution. The "bias" of the estimators is not computed with respect to this districtwide average, however, but with respect to each individual β_i. In other words, our measure of bias is, in fact, the average error (over all precincts) between the "predicted" parameter, say, β_i^b, and its simulated value. For EI, the mean of the posterior distribution is taken as the predicted value (for Goodman's method, the predicted value is, of course, the same in each precinct). Estimates

[25]Technically, the Goodman model *does* provide precinct-level estimates as well. The estimate for each precinct is simply the same as the districtwide estimate. This is the constancy assumption.

of precision are provided by the root mean squared error (RMSE) and the mean absolute error (MAE).[26]

To obtain a benchmark for the discussion of the spatial models that follows, we first consider the results for the base case (reported as the $\rho = 0$ case). These are not difficult cases because only a *very small* degree of heterogeneity and truncation is implied by these parameters.[27] Consequently, these samples represent fairly "stable" cases that are probably more "homogeneous" than those encountered in actual empirical practice.[28] This is confirmed by the "districtwide" true values listed in the second and third columns in Tables 2 and 3. In all cases, these are virtually the same as the mean for the untruncated underlying bivariate normal distribution. For these designs, it is thus not surprising that both EI and OLS perform very similarly in the base case.

It is well known that OLS is unbiased and that the heteroskedasticity affects only the precision of the estimate. In our simulations, since the data were generated to have no aggregation bias, but only spatial effects, it is unsurprising that we find essentially no bias in our Monte Carlo simulations. Hence, we do not report the bias results in the tables. Confirming earlier simulation results, then, we find no bias for either EI or OLS and basically the same precision [see King (1997) for EI results and Cho (1998) for EI and OLS results]. For example, for β_i^b (the truncated model with rook weights), the RMSE is 0.0155 for EI and 0.0154 for OLS; virtually indistinguishable.

Notably, however—and previously overlooked—the RMSE increases as the degree of spatial autocorrelation increases. In Table 2, we can see that for the truncated model (β^b, rook weights), the RMSE grows by a factor of 1.5 as the underlying spatial autocorrelation grows to 0.8. In practice, it is difficult to assess the degree of spatial autocorrelation. Hence, an analyst would be unsure whether the data of interest were in the realm of a low degree of spatial autocorrelation, where the RMSE is not as large, or a high degree of spatial autocorrelation, where the RMSE is larger and points to problems with the precision of the estimator. In addition, in actual practice, one must also consider that spatial autocorrelation is unlikely to occur in isolation (King 1997; Cho 1998). Spatial effects may even signal problems with aggregation bias (Achen and Shively 1995).

Although RMSE includes both a bias component and a precision component, in our case, since there is essentially no bias, RMSE provides practically a pure measure of precision. The results are similar when we examine the MAE values. In our results, then, in terms of precision, we can see that the performance of both EI and OLS begins to deteriorate when the underlying distribution becomes more heterogeneous. It is against this base case that we need to assess the influence of any spatial effects that are introduced into the design.

In some instances, the performance of the EI estimator degrades marginally faster as the degree of spatial autocorrelation increases. In other cases, the performance of the OLS

[26] Note that we do not provide "coverage rates" for, say, an 80% confidence interval. This is a deliberate choice, since our main point is that spatial effects affect the precision of the estimates. Since EI does not correct for spatial autocorrelation, the size of the standard errors is, to us, not obviously useful. Instead, since EI makes no adjustments for the increased precision, its estimates of the standard errors are based on a model that specifies no spatial effects. How useful these uncorrected estimates may be is not a topic we explore here. We focus on the precision of the estimator.

[27] For a normal distribution with a mean of 0.5 and a standard deviation of 0.2, 95% of the variates are contained in the interval 0.1 to 0.9, leaving truncation to only true outliers. The situation is even more pronounced for a standard deviation of 0.1.

[28] Indeed, one may note that in King's one-directional, unidimensional simulations (1997), his results for high levels of "spatial autocorrelation" are in many ways worse than the results we provide here. Our simulations are not intended to, and do not, present the worst possible manner in which spatial effects may occur. As King himself demonstrates, the problems can be more severe.

estimator degrades more quickly as the degree of spatial autocorrelation increases.[29] For none of these spatial models does the performance for either estimator improve or stay the same. In practice, an analyst will have no idea which simulated model bears a closer resemblance to the data of interest, if either spatial model bears any resemblance to his data, or if these spatial effects are confounded by other issues such as aggregation bias. Hence, it is not as important to note the specifics of these models as it is to note that the performance of both estimators deteriorates as the degree of spatial autocorrelation increases. The small differences in performance here are overshadowed by the larger issues that arise in making ecological inferences.

In this context of the base case, one would expect that the effect of increasing the sample size would be an improvement in efficiency (i.e., the RMSE would decrease) for *all* estimators across the board. In his one-directional, unidimensional results, King demonstrates EI's improved performance as the sample size rises from 100 to 1000. He does not revisit this efficiency issue in his subsequent simulations (King 2000). In our analysis of this issue, in terms relative to our base case, when we ran the same simulations for a sample size four times as large, the overall patterns were the same. Hence, we cannot substantiate the claim about efficiency.[30]

Although it is impossible fully to characterize all forms of spatial effects, these few patterns in the performance of the two estimators in the presence of these particular spatial effects are striking. In contrast to earlier evidence regarding the effect of spatial autocorrelation, we find that even in the absence of aggregation bias, there is a separate effect on the precision of the estimators, one that becomes more noticeable as the degree of spatial autocorrelation increases. This result holds despite the highly stylized nature of our data generation process, which is predisposed to produce subtle results. Nonetheless, in our analysis, as the degree of spatial autocorrelation rises (from $\rho = 0$ to $\rho = 0.8$), the RMSE and MAE rise as well. The effect is similar for the rook, queen (not reported), and King spatial weights.[31] Hence, the effect is neither unique to nor driven by the choice of weights. Moreover, because the effect is similar for both the truncated and the censored case, the effect cannot be attributed to either of these model designs.

7.2 *Reconciling Previous Findings*

Our results differ from those reported by King (1997, 2000). We might expect our results to differ from those reported in his book (King 1997), since the data generating process for those simulations, he now agrees, was not indicative of spatial processes (King 2000, p. 603). However, our results differ even from his subsequent simulations (King 2000) when we employ his own weights matrix. He used this particular weights matrix to "ensure a realistic form of multidirectionality and two-dimensionality" (King 2000, p. 603). Upon reporting

[29] The details are available in Tables 2 and 3. In general, the performance of the two estimators is almost the same for the truncated model, with OLS just edging out EI. For the censored model, EI seems to outperform OLS (most notably when $\rho = 0.8$). Nonetheless, the pattern of behavior is the same and both clearly degrade as the degree of spatial autocorrelation increases. It is worth making the point again that these Monte Carlo simulations are stylized. Consequently, the results are necessarily subtle and the pattern of behavior is more noteworthy than the specific RMSE or MAE in each case. As we have stressed previously, spatial effects can take many forms, and we have simulated only two specific cases here. In the only well-documented analysis of realistic spatial effects, then, we find evidence that the precision of both EI and OLS is adversely affected by the presence of spatial effects. *Much* more research must be done to assess the specific impact of all forms of spatial effects.

[30] Results are available at the *Political Analysis* Web site.

[31] Results for the queen weights are available at the *Political Analysis* Web site.

his results, he claimed that his subsequently reported Monte Carlo evidence "confirm[s] conclusions from the simpler analysis presented in [his] book: Spatial autocorrelation has only a minimal effect on model estimates and standard errors" (2000, p. 603). In short, in his Monte Carlo experiments, the average absolute error (MAE) is approximately zero, and the true standard deviation across the simulations is almost precisely the average of the estimated standard errors from each simulation (2000, pp. 603–604).[32] As we have mentioned, our evidence confirms the result regarding bias. However, our findings on precision differ.

It would be helpful to identify the source of our discrepant findings. In this case, however, it is impossible to identify the exact source of the discrepancy, since King's Monte Carlo results are difficult to navigate and interpret based on what he has reported. To be clear, there have been two assessments by King of the role of spatial effects in ecological inference models. His first assessment appears in his 1997 book. That analysis includes generated data that is limited and not representative of spatial processes (Anselin 2000, O'Loughlin 2000). In response, King (2000, p. 603) noted that "[t]hese are reasonable criticisms." He then reports on new simulations that are ostensibly more representative of spatial processes, but the brevity of his report on those simulations makes it impossible for us to fully understand the details of the simulation experiment and thus prohibits us from making meaningful comparisons between his results and ours. Since crucial information is not reported in King (2000) (such as what is the model being estimated, what are the parameters, what is the data generating process?), it is difficult for us to investigate the deviations from his findings and ours. It is clear from the reported results, however, that our Monte Carlo experiments and our assessment, in general, are far more extensive. As just one example, we simulated different degrees of spatial autocorrelation, thus allowing us to make claims about the performance of estimators as this characteristic changes. King (2000) reports just one model with unclear parameters and specification.

In summary, we find that the earlier assessments of the absence of an effect resulting from spatial autocorrelation were too optimistic. Whether introduced in the form of a spatial autoregressive process in the random coefficient specification or as a censoring of an underlying latent random component, the effect exists and is significant. The effect is limited to the precision of the estimator, since these forms of spatial autocorrelation, in and of themselves, do not induce bias. More importantly, the effect is similar for both EI and OLS (Cho 1998). The performance of both estimators declines as the degree of spatial autocorrelation rises.[33]

[32] King's discussion of his model is limited, and this creates some interpretation issues. For instance, he states that "the true standard deviation across the simulations..." In this case, the usage of the word "true" is very confusing, since there are several ways to interpret it. He could be referring to either the "true" variability at the precinct level, the variability of the estimate at the precinct level, or the estimate of the variability of the estimates. We report the variability of the EI estimates around the true value, which are controlled in the simulations. These are not the standard errors computed by EI, which are themselves not "true" values. The only true values are the variances of the true β_i^b and β_i^w, not of the estimated ones. It is unclear that King reports anything related to the true β_i^b and β_i^w that were used to generate the data. Hence, we believe that our results mimic the results that he describes.

[33] One may be tempted to believe that EI must perform better than OLS since OLS produces out-of-bounds estimates. That claim may sound intuitive and simple but is, instead, quite complicated and one that has never been substantiated. While EI always produces in-bounds estimates, these in-bounds estimates are not guaranteed to be correct or even of the right magnitude. Relatedly, while the "constancy assumption" in the OLS model has been widely criticized, EI incoporates an analogous "similarity assumption." That is, while OLS constrains all precincts to have the same parameter values, EI constrains all precincts to have similar (but not identical) parameter values. The close relationship of these two assumptions is the root for much of the similarity between the two estimators.

8 Conclusion

Our results provide additional perspectives on King's (1997, p. 94) assertion that spatial autocorrelation does not pose difficulties for EI. We have identified another assumption (no spatial autocorrelation) that must be met to ensure that EI estimates are problem free. The EI model is affected by violations in this assumption, and the resulting effects when this assumption is violated may pose problems for proper inference. In particular, the presence of a spatial autoregressive process manifests itself as a loss of precision in ecological inference estimators. The effect is thus similar to the standard regression context.

Second, we have demonstrated the impact of considering spatial effects that are much richer than the unidimensional, one-directional time-series analogues examined by King (1997) and the spatial rendition of King (2000). Although we have categorized several effects arising from spatial heterogeneity, note that we have considered only a small subset of the possible interactions between dependence and heterogeneity, and only in the highly stylized situation where the effects of truncation and variance were minimal. However, even in these artificial designs, our results suggest that spatial effects should not be ignored and deserve as central a consideration in the EI context as in standard regression analysis. Hence, there is a need to develop effective tests to assess spatial effects as well as to develop new models that explicitly account for these effects. Importantly, it would seem that any information about the "structure" of the spatial heterogeneity should be exploited before defaulting to the "generic" assumption of random coefficient variation. Certainly, estimation methods that are robust to the presence of various forms of heterogeneity/dependence should be explored.

For the simple designs considered here, our results are in line with earlier findings (Cho 1998) that have demonstrated the virtually identical performance of EI and the "naive" Goodman OLS estimator. Ironically, then, the additional computational complexity of the EI estimator contributes little in the face of spatial effects. These results hold whether the violated assumption is the distributional assumption, the aggregation bias assumption, or the spatial autocorrelation assumption.

We should reiterate that the simulation designs (and those considered by King) are highly stylized and, for one, ignore the interaction between spatial effects and aggregation bias, which may be crucially important in practice. Violations of the spatial autocorrelation assumption may, for instance, signal problems with the aggregation bias assumption. Indeed, it is extremely rare that these violations would not occur simultaneously. Also, we should stress that the bivariate model that follows from the fundamental accounting identity is rather limited for empirical practice in political analysis. While it may be appropriate for Voting Rights court cases, it provides a highly simplified (and too simplistic) model for the interaction between race and political behavior. Once outside the fundamental accounting identity, a much richer set of multivariate models may be considered, for which a wide range of spatial effects can be modeled. We argue that this is a more promising route to pursue in terms of incorporating spatial effects into political models for ecological regression.

References

Achen, Christopher H., and W. Phillips Shively. 1995. *Cross-Level Inference*. Chicago: University of Chicago Press.

Anselin, Luc. 1988. *Spatial Econometrics: Methods and Models*. Boston: Kluwer Academic.

Anselin, Luc. 1990. "Spatial Dependence and Spatial Structural Instability in Applied Regression Analysis." *Journal of Regional Science* 30:185–207.

Anselin, Luc. 1992. "Spatial Dependence and Spatial Heterogeneity: Model Specification Issues in the Spatial Expansion Paradigm." In *Applications of the Expansion Method*, eds. J. P. Jones and E. Casetti. London: Routledge, pp. 334–354.

Anselin, Luc. 1995. "Local Indicators of Spatial Association—LISA." *Geographical Analysis* 27:93–115.

Anselin, Luc. 2000. "The Alchemy of Statistics, or Creating Data Where No Data Exist." *Annals of the Association of American Geographers* 90:93–115.

Anselin, Luc, and A. Bera. 1998. "Spatial Dependence in Linear Regression Models with an Introduction to Spatial Econometrics." In *Handbook of Applied Economic Statistics*, eds. A. Ullah and D. Giles. New York: Marcel Dekker, pp. 237–289.

Ansolabehere, Stephen, and Douglas Rivers. 1997. "Bias in Ecological Regression Estimates." Working paper.

Casetti, Emilio. 1997. "The Expansion Method, Mathematical Modeling, and Spatial Econometrics." *International Regional Science Review* 20:9–33.

Cho, Wendy K. Tam. 1998. "Iff the Assumption Fits...: A Comment on the King Ecological Inference Solution." *Political Analysis* 7:143–163.

Cho, Wendy K. Tam. 2001. "Latent Groups and Cross-Level Inferences." *Electoral Studies* 20:243–263.

Cliff, A., and K. Ord. 1981. *Spatial Processes: Models and Applications*. London: Pion.

Fotheringham, A. Stewart, Chris Brundson, and Charlton Martin. 1998. "Geographically Weighted Regression: A Natural Evolution of the Expansion Method for Spatial Data Analysis." *Environment and Planning A* 30:1905–1927.

Freedman, D. A., S. P. Klein, M. Ostland, and M. R. Roberts. 1998. Book review of "A Solution to the Ecological Inference Problem." *Journal of the American Statistical Association* 93:1518–1522.

Goodman, Leo A. 1953. "Ecological Regressions and Behavior of Individuals." *American Sociological Review* 18:663–664.

Goodman, Leo A. 1959. "Some Alternatives to Ecological Correlation." *American Journal of Sociology* 64:610–625.

Griffiths, William E. 1972. "Estimation of Actual Response Coefficients in the Hildreth-Houck Random Coefficient Model." *Journal of the American Statistical Association* 68:329–335.

Griffiths, William E., Ross G. Drynan, and Surekha Prakash. 1979. "Bayesian Estimation of a Random Coefficient Model." *Journal of Econometrics* 10:201–220.

King, Gary. 1997. *A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data*. Princeton, NJ: Princeton University Press.

King, Gary. 2000. "Geography, Statistics, and Ecological Inference." *Annals of the Association of American Geographers* 90:601–606.

Lancaster, Tony. 2000. "The Incidental Parameter Problem Since 1948." *Journal of Econometrics* 95:391–413.

Neyman, Jerzy, and Elizabeth L. Scott. 1948. "Consistent Estimation from Partially Consistent Observations." *Econometrica* 16:1–32.

O'Loughlin, John. 2000. "Extending King's Ecological Inference Approach to Answer a Social Science Puzzle: Who Voted for the Nazi Party in Weimar Germany?" *Annals of the Association of American Geographers* 90:592–601.

Openshaw, Stan, and Peter Taylor. 1979. "A Million or so Correlation Coefficients: Three Experiments on the Modifiable Areal Unit Problem." In *Statistical Applications in the Spatial Sciences*, ed. N. Wrigley. London: Pion, pp. 127–144.

Robinson, W. S. 1950. "Ecological Correlations and the Behavior of Individuals." *American Sociological Review* 15:351–357.