# Adding Spatially-Correlated Errors Can Mess Up The Fixed Effect You Love

**James S. Hodges**[1*]**, Brian J. Reich**[2]

[1]Division of Biostatistics, University of Minnesota, Minneapolis, Minnesota USA 55414
[2]Department of Statistics, North Carolina State University, Raleigh, North Carolina USA 27695
*email: hodge003@umn.edu

March 4, 2010

## SUMMARY

Many statisticians have had the experience of fitting a linear model with uncorrelated errors, then adding a spatially-correlated error term (random effect) and finding that the estimates of the fixed-effect coefficients have changed substantially. We show that adding a spatially-correlated random effect to a linear model is equivalent to adding a saturated collection of canonical regressors, the coefficients of which are shrunk toward zero, where the spatial map determines both the canonical regressors and the relative extent of the coefficients' shrinkage. Adding a spatially-correlated random effect can also be seen as inflating the error variances associated with specific contrasts of the data, where the spatial map determines the contrasts and the extent of error-variance inflation. We show how to restrict the spatial random effect to the orthogonal complement of the fixed effects, which we call restricted spatial regression. We mostly model spatial correlation using an improper conditional auto-regression (ICAR), but briefly show that spatial confounding also arises with so-called geostatistical models and penalized splines, and for the same reason as with the ICAR. We consider five proposed interpretations of spatial confounding and draw implications about what, if anything, one should do about it. For a given problem, the appropriate action depends on whether the spatial random effect is merely as a formal device used to implement spatial smoothing, or a random effect in the traditional sense of, say, Scheffé (1959). For spatial random effects with the former interpretation, restricted spatial regression should be used, while for the latter interpretation this is less clear. In the process, we debunk the common belief that adding a spatially-correlated random effect adjusts fixed effect estimates for spatially-structured missing covariates.

Key Words: confounding, missing covariate, random effect, spatial correlation, spatial regression

# 1 Stomach cancer in Slovenia: Where does the fixed effect go?

Dr. Vesna Zadnik, a Slovenian epidemiologist, collected data describing counts of stomach cancers in the 194 municipalities that partition Slovenia, for the years 1995 to 2001 inclusive. She was studying the possible association of stomach cancer with socioeconomic status, as measured by a composite score calculated from 1999 data by Slovenia's Institute of Macroeconomic Analysis and Development. (Her findings were published in Zadnik & Reich 2006.) Figure 1a shows the standardized incidence ratio ($SIR$) of stomach cancer for the 194 municipalities; for municipality $i = 1, \ldots, 194$, $SIR_i = O_i/E_i$, where $O_i$ is the observed count of stomach cancer cases and $E_i$ is the expected count using indirect standardization. Figure 1b shows the socioeconomic scores for the municipalities, $SEc_i$, after centering and scaling, so the $SEc_i$ have average 0 and finite-sample variance 1. In both panels of Figure 1, dark colors indicate larger values. $SIR$ and $SEc$ have a negative association: western municipalities generally have low $SIR$ and high $SEc$ while eastern municipalities generally have high $SIR$ and low $SEc$.

Following advice received in a spatial-statistics short course, Dr. Zadnik first did a non-spatial analysis assuming the $O_i$ were independent Poisson observations with $\log\{E(O_i)\} = log(E_i) + \alpha + \beta SEc$, with flat priors on $\alpha$ and $\beta$. This analysis gave the obvious result: $\beta$ had posterior median -0.14 and 95% posterior interval $(-0.17, -0.10)$, capturing Figure 1's negative association.

Dr. Zadnik continued following the short course's guidance by doing a spatial analysis using the improper (or implicit) conditionally autoregressive (ICAR) model of Besag, York, & Mollié (1991). Dr. Zadnik's understanding was that ignoring spatial correlation would make $\beta$'s posterior standard deviation (standard error) too small, while spatial analysis in effect discounts the sample size with little effect on the estimate of $\beta$, just as generalized estimating equations (GEE) adjusts standard errors for clustering but (in the authors' experience) has little effect on point estimates unless the working correlations are very large. As we will see, other people have different reasons for introducing spatial correlation.

In the model of Besag, York, & Mollié (1991), the $O_i$ are conditionally independent Poisson random variables with mean

$$\log\{E(O_i)\} = log(E_i) + \beta SEc + S_i + H_i. \tag{1}$$

The intercept is now the sum of two random effects, $\mathbf{S} = (S_1, \ldots, S_{194})'$ capturing spatial clustering and $\mathbf{H} = (H_1, \ldots, H_{194})'$ capturing heterogeneity. The $H_i$ are modeled as independent draws from a normal distribution with mean zero and precision (reciprocal of variance) $\tau_h$. The $S_i$ are modeled using an $L_2$-norm ICAR, also called a Gaussian Markov random field, which is discussed in detail below. The ICAR represents the intuition that neighboring municipalities tend to be more similar to each other than municipalities that are far apart, where similarity of neighbors is controlled by an unknown parameter $\tau_s$ that is like a precision. This Bayesian analysis used independent gamma priors for $\tau_h$ and $\tau_s$ with mean 1 and variance 100, and a flat prior for $\beta$.

In the spatial analysis, $\beta$ had posterior mean -0.02 and 95% posterior interval $(-0.10, 0.06)$. Compared to the non-spatial analysis, the 95% interval was wider and the spatial model fit better,

with deviance information criterion (DIC; Spiegelhalter et al 2002) decreasing from 1153 to 1082 even though the effective number of parameters ($p_D$) increased sharply, from 2.0 to 62.3. These changes were expected. The surprise was that the negative association, which is quite plain in the maps and the non-spatial analysis, had disappeared. What happened?

This spatial confounding effect, which we reported in Reich, Hodges, & Zadnik (2006), has been reported elsewhere but is not widely known. The earliest report we have found is Clayton, Bernardinelli, & Montomoli (1993), who used the term "confounding due to location" for a less dramatic but still striking effect in analyses of lung-cancer incidence in Sardinia; they and Wakefield (2007) report a similar-sized effect in the long-suffering Scottish lip-cancer data. This effect is not yet understood; we have seen five proposed interpretations. The interpretations depend on whether the random effect **S** meets the traditional definition of "random effect" given in, for example, Scheffé (1959, p. 238): the levels of a random effect are draws from a population, and the draws are not of interest in themselves but only as samples from the larger population, which *is* of interest. For reasons to be discussed later, we describe random effects *not* meeting this definition as formal devices to implement a smoother. With this distinction, which Section 3 discusses in more detail, the five interpretations are:

- **The random effect S is a formal device to implement a smoother.**

    - (i) Spatially-correlated errors remove bias in estimating $\beta$ and are generally conservative (Clayton et al 1993).

    - (ii) Spatially-correlated errors can introduce or remove bias in estimating $\beta$ and are not necessarily conservative (Wakefield 2007; implicit in Reich et al 2006).

- **S is a Scheffé-style random effect.**

    - (iii) The spatial effect **S** is collinear with the fixed effect, but neither estimate of $\beta$ is biased (David B. Nelson, personal communication).

    - (iv) Adding the spatial effect **S** creates information loss, but neither estimate of $\beta$ is biased (David B. Nelson, personal communication).

    - (v) Because error is correlated with the regressor $SEc$ in the sense commonly used in econometrics, *both* estimates of $\beta$ are biased (Paciorek 2009).

Except for (v), these interpretations treat $SEc$ as measured without error and not drawn from a probability distribution.

Our purpose is to examine these interpretations and determine which is appropriate under what circumstances. Section 2 describes the mechanics of spatial confounding. We give derivations for the normal-errors analog to (2); Reich et al (2006) gave some of these derivations and an extension to generalized linear mixed models. Briefly, adding the ICAR-modeled effect **S** to the model is equivalent to adding a saturated collection of canonical regressors that are determined solely by the spatial map. The coefficients of these canonical regressors are shrunk toward zero; their relative

degrees of shrinkage are also determined solely by the spatial map. Adding the random effect $\mathbf{S}$ can also be seen as inflating the error variances associated with specific contrasts of the data, where the spatial map determines the contrasts and the extent of error-variance inflation. These derivations suggest a method to restrict spatial smoothing to the orthogonal complement of the fixed effects, $SEc$ for the Slovenian data, which we call restricted spatial regression. Spatial confounding is also present for so-called geostatistical models and two-dimensional penalized splines which, like ICAR, represent the intuition that observations taken near each other in space tend to be more similar than observations taken far apart in space. Section 3 considers the five interpretations in light of Section 2's theory and draws implications about what, if anything, one should do about spatial confounding. For a given problem, the appropriate action depends on whether the spatial random effect is interpreted merely as a formal device used to implement spatial smoothing. If so, restricted spatial regression should be the preferred analysis. In making this argument, we debunk the common belief that introducing a spatially-correlated random effect adjusts fixed effect estimates for spatially-structured missing covariates (e.g., Clayton et al 1993) and argue in addition that from the linear-models perspective adopted throughout, it seems perverse to even try to use an error term to adjust a coefficient estimate. For Scheffé-style random effects, on the other hand, the preferred action is as yet unclear.

# 2 The mechanics of spatial confounding

This section relies heavily on the theory of linear models. Although our approach is Bayesian, most of our results apply immediately to analyses based on maximizing the restricted likelihood.

## 2.1 The model with spatial correlation written as a linear model

For an $n$-dimensional observation vector $\mathbf{y}$, write the normal-errors analog to model (1) as

$$\mathbf{y} = \mathbf{X}\beta + \mathbf{I}_n\mathbf{S} + \boldsymbol{\epsilon} \tag{2}$$

where $\mathbf{y}$, $\mathbf{X}$, $\mathbf{S}$, and $\boldsymbol{\epsilon}$ are $n \times 1$, $\beta$ is scalar, $\mathbf{I}_n$ is the $n$-dimensional identity matrix, $\mathbf{y}$ and $\mathbf{X}$ are known, and $\beta$, $\mathbf{S}$, and $\boldsymbol{\epsilon}$ are unknown. In the Slovenian data, $n = 194$ and $\mathbf{X} = SEc$. $\mathbf{X}$ is centered and scaled to have average 0 and finite-sample variance 1. The derivation below generalizes easily to any full-rank $\mathbf{X}$ without an intercept column (Reich et al 2006); the intercept is implicit in $\mathbf{S}$, as shown below. The error term $\boldsymbol{\epsilon}$ is $n$-dimensional normal with mean zero and precision matrix $\tau_e\mathbf{I}$, $\tau_e$ being the reciprocal of the error variance. Our Bayesian analysis puts a flat prior on $\beta$, but this is not necessary.

The $L_2$-norm improper CAR model (or prior) on $\mathbf{S}$ can be represented as an improper $n$-variate normal distribution specified by its precision matrix:

$$p(\mathbf{S}|\tau_s) \propto \tau_s^{(n-G)/2} \exp(-0.5\tau_s\mathbf{S}'\mathbf{Q}\mathbf{S}) \tag{3}$$

where $G$ is the number of islands (disconnected groups of municipalities) in the spatial map (Hodges, Carlin, & Fan 2003). The unknown $\tau_s$ controls the smoothness of $\mathbf{S}$; larger $\tau_s$ force neighboring

4

$S_i$ to be more similar to each other. $\mathbf{Q}$ encodes the neighbor pairs, with diagonal elements $q_{ii} =$ number of municipality $i$'s neighbors, and $q_{ij} =$ -1 if municipalities $i$ and $j$ are neighbors and 0 otherwise. For the Slovenian data, $G = 1$ and we specify that municipalities $i$ and $j$ are neighbors if they share a boundary. Other specifications of neighbor pairs are possible but this one is common. Model (3) can be re-expressed in the pairwise-difference form

$$p(\mathbf{S}|\tau_s) \propto \tau_s^{(n-G)/2} \exp(-0.5\tau_s \sum (S_i - S_j)^2) \tag{4}$$

where the sum is over unique neighbor pairs $(i, j)$.

We have written model (2) with the random effect $\mathbf{S}$ having design matrix $\mathbf{I}_n$ to emphasize that model (2) is over-parameterized and is identified only because $\mathbf{S}$ is smoothed or, alternatively, constrained by the ICAR model (3). To clarify the identification issues, re-parameterize (2) as follows. The neighbor-pair matrix $\mathbf{Q}$ has spectral decomposition $\mathbf{Q} = \mathbf{ZDZ}'$, where $\mathbf{Z}$ is $n \times n$ orthogonal and $\mathbf{D}$ is diagonal with diagonal elements $d_1 \geq \ldots \geq d_{n-G} > 0$ and $d_{n-G+1} = \ldots = d_n = 0$ (Hodges et al 2003). $\mathbf{Z}$'s columns are $\mathbf{Q}$'s eigenvectors; $\mathbf{D}$'s diagonal elements are the corresponding eigenvalues. Re-parameterize (2) as

$$\mathbf{y} = \mathbf{X}\beta + \mathbf{Zb} + \epsilon, \tag{5}$$

where $\mathbf{b} = \mathbf{Z}'\mathbf{S}$ is an $n$-vector with a normal distribution having mean zero and diagonal precision matrix $\tau_s\mathbf{D} = \tau_s diag(d_1, \ldots, d_{n-G}, 0, \ldots, 0)$.

The spatial random effect $\mathbf{S}$ thus corresponds to a saturated collection of canonical regressors, the $n$ columns of $\mathbf{Z}$, whose coefficients $\mathbf{b}$ are shrunk toward zero to an extent determined by $\tau_s\mathbf{D}$. The smoothing parameter $\tau_s$ controls shrinkage of all $n$ components of $\mathbf{b}$ and the $d_i$ control the relative degrees of shrinkage of the $b_i$ for a given $\tau_s$. Both the canonical regressors $\mathbf{Z}$ and the $d_i$ are determined solely by the spatial map through $\mathbf{Q}$. The first canonical coefficient, $b_1$, has the largest $d_i$ and is thus shrunk the most for any given $\tau_s$; for the Slovenian map, $d_1 = 14.46$. The last $G$ $b_i$, $b_{n-G+1}, \ldots, b_n$ are not shrunk at all because their prior precisions $\tau_s d_i$ are zero, so they are fixed effects implicit in the spatial random effect. The $b_i$ with the smallest positive $d_i$, $d_{n-G}$, is shrunk least of all the shrunken coefficients. For the Slovenian map with $G = 1$, this is $b_{193}$, with $d_{193} = 0.03$, so its prior precision is smaller than $b_1$'s by a factor of about 500 for any $\tau_s$.

To understand the differential shrinkage of the $b_i$, we need to understand the columns of $\mathbf{Z}$, which the spatial map determines. $Z_{n-G+1}, \ldots, Z_n$, whose coefficients $b_{n-G+1}, \ldots, b_n$ are not shrunk at all, span the space of the means of (or intercepts for) the $G$ islands in the spatial map. This is easily seen from (4): the improper CAR distribution is flat (puts no constraint) on the $G$ island means. Also, $\mathbf{Q1}_n = 0$ for any map, so the overall intercept $\mathbf{1}_n$ always lies in the span of $Z_{n-G+1}, \ldots, Z_n$ and is thus implicit in the ICAR specification. Thus without loss of generality, we set $Z_n = \frac{1}{\sqrt{n}}\mathbf{1}_n$ so all other $Z_i$ are contrasts, i.e., $\mathbf{1}_n'Z_i = 0$.

Based on examples, $Z_{n-G}$, whose coefficient $b_{n-G}$ has the smallest positive prior precision $\tau_s d_{n-G}$, can be interpreted loosely as the lowest *frequency* contrast in $\mathbf{S}$ among the shrunken contrasts. Figure 2a is a plot of $Z_{n-1} = Z_{193}$ for the Slovenian data, where darker and lighter colors indicate higher and lower values respectively. $Z_{193}$ is "low-frequency" in the sense that it

is a roughly linear trend along the long axis of Slovenia's map. Again based on examples, as the value of $d_i$ increases, the frequency of the corresponding $Z_i$ increases as well. Figure 2b shows $Z_1$ for the Slovenian data; it is roughly the difference between the two municipalities with the most neighbors (the dark municipalities) and the average of their neighbors. For other spatial maps, the interpretations are similar. For the counties of Minnesota (Reich & Hodges 2008, Fig. 6), the contrast with the least-smoothed coefficient is the north-south gradient, while the contrast with the most-smoothed coefficient is roughly the difference between the county having the most neighbors and the average of those neighbors. For the spatial map representing periodontal measurement sites, the contrast with the least-smoothed coefficient is nearly linear along a dental arch, while the contrasts with the most-smoothed coefficients are the difference, on each tooth, between the average of the interproximal sites and the average of the direct sites (Reich & Hodges 2008, Fig. 3, which also shows other $Z_i$). For the Scottish lip-cancer data, the contrast with the least-smoothed coefficient is the north-south gradient, while the contrast with the second least-smoothed coefficient is roughly quadratic along the north-south gradient, high at the northern and southern extremes and low in the middle — just like the much-studied predictor of lip cancer, $AFF$, which measures employment in agriculture, fisheries, and forestry.

## 2.2 Spatial confounding explained in linear-model terms

For the normal-errors model (2), it is straightforward to show (Reich et al 2006) that the posterior mean of $\beta$ conditional on the precisions $\tau_e$ and $\tau_s$ is

$$E(\beta|\tau_e, \tau_s, \mathbf{y}) = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Z}E(\mathbf{b}|\tau_e, \tau_s, \mathbf{y}), \tag{6}$$

where $E(\mathbf{b}|\tau_e, \tau_s, \mathbf{y})$ is *not* conditional on $\beta$, taking the value

$$E(\mathbf{b}|\tau_e, \tau_s, \mathbf{y}) = (\mathbf{Z}'\mathbf{P}^c\mathbf{Z} + r\mathbf{D})^{-1}\mathbf{Z}'\mathbf{P}^c\mathbf{y}, \tag{7}$$

for $r = \tau_s/\tau_e$ and $\mathbf{P}^c = \mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$, the familiar residual projection matrix from linear models. These expressions are correct for full-rank $\mathbf{X}$ of any dimension.

In (6), the term $(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$ is the ordinary least-squares estimate of $\beta$ and also $\beta$'s posterior mean in a fit without $\mathbf{S}$, using a flat prior on $\beta$. Thus the second term $-(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Z}E(\mathbf{b}|\tau_e, \tau_s, \mathbf{y})$ is the change in $\beta$'s posterior mean, conditional on $(\tau_e, \tau_s)$, from adding $\mathbf{S}$. Note that $\mathbf{Z}E(\mathbf{b}|\tau_e, \tau_s, \mathbf{y})$ is the fitted value of $\mathbf{S}$ given $(\tau_e, \tau_s)$. Thus when $\mathbf{S}$ is added to the model, the change in $\beta$'s posterior mean given $(\tau_e, \tau_s)$ is -1 times the regression on $\mathbf{X}$ of the fitted values of $\mathbf{S}$.

Because $\mathbf{X}$ is centered and scaled (so $\mathbf{X}'\mathbf{X} = n - 1$) and $\mathbf{Z}$ is orthogonal, the correlations $\rho_i$ of $\mathbf{X}$ and $Z_i$, the $i^{th}$ column of $\mathbf{Z}$, can be written as $\mathbf{R} = (\rho_1, \ldots, \rho_{n-G}, 0, \ldots, 0)' = (n-1)^{-0.5}\mathbf{X}'\mathbf{Z}$. Equation (6) can then be written as

$$E(\beta|\tau_e, \tau_s, \mathbf{y}) = \hat{\beta}_{OLS} - (n-1)^{1/2}\mathbf{R}'E(\mathbf{b}|\tau_e, \tau_s, \mathbf{y}). \tag{8}$$

(Appendix C gives a more explicit expression.) From (8), adding $\mathbf{S}$ to the model induces a large change in $\beta$'s posterior mean if $\rho_i$ and $E(b_i|\tau_e, \tau_s, \mathbf{y})$ are large for the same $i$. This happens if four

6

conditions hold: $\mathbf{X}$ is highly correlated with $Z_i$; $\mathbf{y}$ has a substantial correlation with both $\mathbf{X}$ and $Z_i$; $r = \tau_s/\tau_e$ is small; and $d_i$ is small. The last two conditions ensure that $b_i$ is not shrunk much toward zero. These necessary conditions are all present in the Slovenian data: Figure 1 shows the strong association of $\mathbf{y}$ and $\mathbf{X} = SEc$, $\rho_{193} = correlation(SEc, Z_{193}) = 0.72$, $d_{193} = 0.03$, and $\mathbf{S}$ is not smoothed much (the effective number of parameters in the fit is $p_D = 62.3$).

This effect on $\beta$'s estimate is easily understood in linear-model terms as the effect of adding a collinear regressor to a linear model. If $\mathbf{S}$ were not smoothed at all — if the coefficients $\mathbf{b}$ of the saturated design matrix $\mathbf{Z}$ were not shrunk toward zero — then $\beta$ would not be identified. This corresponds to setting the smoothing precision $\tau_s$ to 0, so the smoothing ratio $r = \tau_s/\tau_e = 0$. If the smoothing ratio $r$ is small, $\beta$ is identified but the coefficients of $Z_i$ with small $d_i$ are shrunk very little, so if these $Z_i$ are collinear with $\mathbf{X}$, the estimate of $\beta$ is subject to the same collinearity effects as in linear models.

Consider also how $\beta$'s posterior variance changes when the ICAR-distributed $\mathbf{S}$ is added to the model. Reich et al (2006) showed that conditional on $\tau_e$ and $\tau_s$, adding $\mathbf{S}$ to the model multiplies the conditional posterior variance of $\beta$ by

$$\left[ 1 - \sum \frac{\rho_i^2}{1 + rd_i} \right]^{-1}, \tag{9}$$

where the sum is over $i$ with $d_i > 0$ and $\rho_i$ is as above. Expression (9) holds only for single-column $\mathbf{X}$; Reich et al (2006) give an expression for general full-rank $\mathbf{X}$. In general, $\sum_{i=1}^{n} \rho_i^2 = 1$. Because $Z_n \propto \mathbf{1}_n$ and $\mathbf{X}$ is centered, $\rho_n = (n-1)^{-0.5} \mathbf{X}' Z_n = 0$, so if $r = 0$, i.e., $\mathbf{S}$ is not smoothed, $var(\beta | \tau_e, \tau_s, \mathbf{y})$ is infinite because $\sum \rho_i^2/(1 + rd_i) = \sum_{i=1}^{n-1} \rho_i^2 = 1$. As $r$ grows from zero, $\mathbf{S}$ is smoothed more and $var(\beta | \tau_e, \tau_s, \mathbf{y})$ decreases. For given $r$, the variance inflation factor (9) is large if $\rho_i$ is large for the smallest $d_i$, as in the Slovenian data. Again, this differs from the analogous result in linear-model theory only because the $b_i$ are shrunk.

## 2.3 Spatial confounding explained in a more spatial-statistics style

The linear-models explanation above seems odd to many spatial-statistics mavens, who usually think of the CAR as a model for the error covariance. We now give a derivation more in line with this viewpoint.

Begin with the re-parameterized normal-errors model (5) but rewrite it in a more spatial-statistics style as $\mathbf{y} = \mathbf{X}\beta + \psi$, where the elements of $\psi = \mathbf{Zb} + \boldsymbol{\epsilon}$ are not independent. $\mathbf{S}$ does not have a proper covariance matrix under an ICAR model, so we must proceed indirectly. Partition $\mathbf{Z} = (\mathbf{Z}^{(1)}|\mathbf{Z}^{(2)})$, where $\mathbf{Z}^{(1)}$ has $n-G$ columns and $\mathbf{Z}^{(2)}$ has $G$ columns, and partition $\mathbf{b}$ conformably as $\mathbf{b} = (\mathbf{b}^{(1)}|\mathbf{b}^{(2)})'$, so $\mathbf{b}^{(1)}$ is $(n-G) \times 1$ and $\mathbf{b}^{(2)}$ is $G \times 1$. Pre-multiply (5) by $\mathbf{Z}'$, so (5) becomes

$$\begin{aligned}
\mathbf{Z}^{(1)'}\mathbf{y} &= \mathbf{Z}^{(1)'}\mathbf{X}\beta + \mathbf{e}_1, & \text{precision}(\mathbf{e}_{1i}) = \tau_e(rd_i)/(1 + rd_i) < \tau_e \\
\mathbf{Z}^{(2)'}\mathbf{y} &= \mathbf{Z}^{(2)'}\mathbf{X}\beta + \mathbf{b}^{(2)} + \mathbf{e}_2, & \text{precision}(\mathbf{e}_{2i}) = \tau_e
\end{aligned} \tag{10}$$

where $\mathbf{e}_1 = \mathbf{b}^{(1)} + \mathbf{Z}^{(1)'}\boldsymbol{\epsilon}$ and $\mathbf{e}_2 = \mathbf{Z}^{(2)'}\boldsymbol{\epsilon}$, so the two rows of (10) are independent, and recall $\tau_e$ is $\boldsymbol{\epsilon}$'s error precision in (5).

Suppose $G = 1$ as in the Slovenian data; $G > 1$ is discussed below. Then $\mathbf{Z}^{(2)} \propto \mathbf{1}_n$ and $\mathbf{Z}^{(2)'}\mathbf{X} = 0$ because $\mathbf{X}$ is centered, so all the information about $\beta$ comes from the first row of (10). Fix $\tau_s$ and $\tau_e$. Without $\mathbf{S}$ in the model, $\beta$'s posterior precision (the reciprocal of $\beta$'s posterior variance) is $\mathbf{X}'\mathbf{X}\tau_e = (n-1)\tau_e$, which is also $\beta$'s information matrix (a scalar, in this case). By adding $\mathbf{S}$ to the model, the posterior precision of $\beta$ decreases by $(n-1)\tau_e \sum_{i=1}^{n-1} \rho_i^2/(1 + rd_i)$, where as before $\rho_i = correlation(\mathbf{X}, Z_i) = (n-1)^{-1/2}\mathbf{X}'Z_i$ and $r = \tau_s/\tau_e$. The information loss is large if $r$ is small (relatively little spatial smoothing) and $\rho_i$ is large for $i$ with small $d_i$, as in the Slovenian data. Because the information about $\beta$ in the different rows of (10) may not be entirely consistent, if row $i$ of (10) differs from the other rows and row $i$ is effectively deleted by the combination of large $\rho_i$ and small $d_i$, $\beta$'s estimate can change markedly when $\mathbf{S}$ is added to the model.

If the spatial map has $G > 1$ islands, the ICAR model includes an implicit $G - 1$ degree of freedom unsmoothed fixed effect for the island means, in addition to the overall intercept. This island-mean fixed effect may be collinear with $\mathbf{X}$ in the usual manner of linear models even if $\mathbf{S}$ is smoothed maximally within islands (i.e., $\tau_s$ is very large).

## 2.4   Avoiding spatial confounding: Restricted spatial regression

Spatial confounding can be interpreted in linear-model terms as a collinearity problem. This suggests adapting a simple trick used as a remedy for collinearity in linear models.

In statistical practice, one is commonly asked to determine whether two variables $\mathbf{X}_1$ and $\mathbf{X}_2$ are independently associated with a third variable $\mathbf{y}$. This vague question is often formalized as a linear model with $\mathbf{y}$ as the dependent variable and $\mathbf{X}_1$ and $\mathbf{X}_2$ as independent variables. If $\mathbf{X}_1$ and $\mathbf{X}_2$ are highly correlated with each other, it can happen that neither is anywhere near statistically significant when both are included in the linear model, even though each is highly significant when it alone is included. Many regression classes teach the following trick for this situation: To test whether $\mathbf{X}_1$ is independently associated with $\mathbf{y}$ in the presence of $\mathbf{X}_2$, regress $\mathbf{y}$ on $\mathbf{X}_2$ and $\mathbf{X}_1^* = (\mathbf{I} - \mathbf{X}_2(\mathbf{X}_2'\mathbf{X}_2)^{-1}\mathbf{X}_2')\mathbf{X}_1$ and test whether $\mathbf{X}_1^*$'s coefficient is zero. $\mathbf{X}_1^*$ is the residuals of $\mathbf{X}_1$ regressed on $\mathbf{X}_2$, so $\mathbf{X}_1^*$ and $\mathbf{X}_2$ are orthogonal by construction and their coefficients' standard errors are therefore not inflated by collinearity.

The analog in our spatial problem is to restrict the spatial random effect $\mathbf{S}$ to the subspace of $n$-dimensional space orthogonal to the fixed effect $\mathbf{X}$, which we call "restricted spatial regression". We show how to do this for a 1-dimensional $\mathbf{X}$, which is easily generalized to higher dimensions (Reich et al 2006, Sec. 3; Sec. 4 extends the method to non-normal observables).

The simplest way to specify a restricted spatial regression is to replace model (2) with $\mathbf{y} = \mathbf{X}\beta + \mathbf{P}^c\mathbf{S} + \epsilon$. The design matrix in front of $\mathbf{S}$ has changed from $\mathbf{I}_n$ to $\mathbf{P}^c = \mathbf{I}_n - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$, the residual projection matrix for a regression on $\mathbf{X}$, but otherwise the model is unchanged. Written this way, $\mathbf{S}$ has a superfluous dimension: one linear combination of $\mathbf{S}$, $(\mathbf{I}_n - \mathbf{P}^c)\mathbf{S}$, necessarily contributes nothing to the fitted values of $\mathbf{y}$ and the data provide no information about it.

For the spatial models considered here, it is easy to reformulate the restricted spatial regression so it has no superfluous dimensions. Let $\mathbf{P}^c$ have spectral decomposition $\mathbf{P}^c = (\mathbf{L}|\mathbf{K})\mathbf{\Phi}(\mathbf{L}|\mathbf{K})'$,

where $\mathbf{\Phi}$ is a diagonal matrix with $n-1$ eigenvalues of 1 and one 0 eigenvalue, $\mathbf{L}$ has $n$ rows and $n-1$ columns, and $\mathbf{K}$ has $n$ rows and 1 column, with $\mathbf{K}$ proportional to $\mathbf{X}$ and $\mathbf{K}'\mathbf{L} = 0$. Then fit the following model:

$$\mathbf{y} = \mathbf{X}\beta + \mathbf{L}\mathbf{S}^* + \boldsymbol{\epsilon}, \tag{11}$$

where $\mathbf{S}^*$ is $(n-1)$-dimensional normal with mean 0 and precision matrix $\tau_s \mathbf{L}'\mathbf{Q}\mathbf{L}$, $\mathbf{Q}$ is the neighbor matrix from $\mathbf{S}$'s ICAR model, and $\boldsymbol{\epsilon}$ is iid normal with mean 0 and precision $\tau_e$.

Using this model in either form and conditioning on $(\tau_s, \tau_e)$, $\beta$ has the same conditional posterior mean as in the analysis without the ICAR-distributed $\mathbf{S}$, but has larger conditional posterior variance (Reich et al 2006, Sec. 3). Thus, restricted spatial regression discounts the sample size to account for spatial correlation without changing $\beta$'s point estimate conditional on $\tau_s$ and $\tau_e$.

## 2.5   Spatial confounding is not an artifact of the ICAR model

Spatial confounding also occurs if, in model (2), $\mathbf{S}$ is given a proper multivariate normal distribution with any of several covariance matrices capturing the intuition that near regions are more similar than distant. Wakefield (2007), for example, found very similar spatial-confounding effects in the Scottish lip-cancer data using the ICAR model and a so-called geostatistical model. For the Slovenian data, we considered geostatistical models in which each municipality's $X_i$ and $y_i$ were treated as being measured at one point in the municipality. For the analyses described below and in Appendix A, each municipality's point had east-west coordinate the average of the farthest east and farthest west boundary points, and analogous north-south coordinate. An example of a proper covariance matrix is $\mathrm{cov}(\mathbf{S}) = \sigma_s^2 \exp(-\delta_{ij}/\theta)$, with $\delta_{ij}$ being Euclidean distance between the points representing municipalities $i$ and $j$ and $\theta$ controlling spatial correlation. Each such covariance matrix that we considered had an unknown parameter like $\theta$, which for now we treat as fixed and known. Applying Section 2.2's approach to such models for $\mathbf{S}$ requires only one change, arising because the precision matrix $\mathrm{cov}(\mathbf{S})^{-1}$ now has no zero eigenvalues: we must add an explicit intercept to model (5). Therefore, holding $\theta$ fixed, spatial confounding will occur by the same mechanism as with the ICAR model for $\mathbf{S}$. (For the Slovenian data, the smallest eigenvalue of $\mathrm{cov}(\mathbf{S})^{-1}$ has eigenvector $Z_{194}$ which, while not constant over the map as for the ICAR model, nonetheless varies little. We have seen this in other geostatistical models; it explains why the intercept can be poorly identified in such models, but the generality of this intercept confounding is unknown.)

In the preceding paragraph, we treated as fixed the parameter $\theta$ in $\mathrm{cov}(\mathbf{S})$. In a non-Bayesian analysis these parameters are typically estimated, while in a Bayesian analysis they are random variables. Thus in applying Section 2.2's analysis to geostatistical models for $\mathbf{S}$, the canonical regressors $\mathbf{Z}$ depend on the unknown parameter $\theta$ of $\mathrm{cov}(\mathbf{S})$. However, for at least four common forms of $\mathrm{cov}(\mathbf{S})$ that we explored (Appendix A), as $\theta$ varied over a wide range the eigenvector corresponding to the second-smallest eigenvalue of $\mathrm{cov}(\mathbf{S})^{-1}$ — the canonical regressor $Z_i$ most likely to produce spatial confounding — was highly correlated with the analogous eigenvector of the ICAR model's precision matrix $\tau_s\mathbf{Q}$, the canonical regressor that *did* produce spatial confounding in

this dataset. Thus, spatial confounding will occur in the Slovenian data for these four geostatistical models.

Applying the spectral approximation to geostatistical models (Appendix B) is another way to make the foregoing more concrete. For measurements taken on a regular square grid, the spectral approximation is closely analogous to Section 2.1's formulation of the ICAR model as a linear model. For this special case, the canonical regressors with least- and most-shrunk coefficients are literally trigonometric functions that are low- and high-frequency respectively.

Penalized splines are a quite different approach to spatial smoothing; Ruppert et al (2003, ch. 13) give a very accessible introduction. However, penalized splines produce the same confounding effect in the Slovenian data. In a class project (Salkowski 2008), a student fit a two-dimensional penalized spline to the Slovenian data, attributing each municipality's counts to a point as described above. He used the R package SemiPar (version 1.0-2; the package and documentation are at http://www.uow.edu.au/∼mwand/webspr/rsplus.html) to fit a model in which municipality $i$'s observed count of stomach cancers $O_i$ was Poisson with log mean

$$\log\{E_i\} + \beta_0 + \beta_{SEc}SEc_i + \beta_E E_i + \beta_N N_i + \sum_k u_k basis_{ki}, \tag{12}$$

where $E_i$ is municipality $i$'s east-west coordinate, $N_i$ is its north-south coordinate (each coordinate was centered and both were scaled by a single scaling constant to preserve the map's shape), $basis_{ki}$ is the default basis in SemiPar (based on the Matérn covariance function), the $u_k$ were modeled as iid normal, and the knots were SemiPar's default knots. SemiPar's default fitting method, penalized quasi-likelihood, shrank the random effect term $\sum_k u_k basis_{ki}$ to zero but the two fixed effects implicit in the spline, $E_i$ and $N_i$, remained in the model and produced a collinearity effect as in an ordinary linear model. Without the spatial spline, a simple generalized linear model fit gave an estimate for $\beta_{SEc}$ of -0.137 (standard error 0.020), essentially the same as in Zadnik's Bayesian analysis, while adding just the fixed effects $E_i$ and $N_i$ changed $\beta_{SEc}$'s estimate to -0.052 (SE 0.028). As the spline fit was forced to be progressively less smooth, $\beta_{SEc}$'s estimate increased monotonically and eventually became positive. (Steinberg & Bursztyn 2004, p. 415, note in passing a similar confounding effect in a different spline.)

Thus, spatial confounding is not an artifact of the ICAR model, but arises from other, perhaps all, specifications of the intuition that measures taken at locations near to each other are more similar than measures taken at distant locations.

# 3 Evaluating the five interpretations; implications for practice

Before we can discuss interpretations of spatial confounding, we need to distinguish two interpretations of "random effect". One is traditional, as in the definition from Scheffé (1959) noted above: the levels of a random effect are draws from a population, and the draws are not of interest in themselves but only as samples from the larger population, which *is* of interest. In recent years,

"random effect" has come to be used in a second sense, to describe effects that have the mathematical form of a Scheffé-style random effect but which are quite different. For these newer-style random effects, the levels are the entire population; or the levels are themselves of interest; or the levels are in no meaningful sense draws from a population, from which further draws could be made. The Slovenian data is an example in which the levels (municipalities) are the entire population. Hospital quality-of-care studies provide examples in which the levels (hospitals) may be considered draws from a population but are themselves of interest. The mixed-model representation of penalized splines (Ruppert et al, 2003) is an example of random effects with levels that are not draws from any conceivable population. In the simplest case of a one-dimensional penalized spline with a truncated-line basis, the random effect is the changes in the fit's slope at each knot, and its distribution is simply a device for penalizing changes in the slope. The senselessness of imagining further draws from such random effects is clearest for the examples in Ruppert et al (2003) in which penalized splines are used to estimate smooth functions in the physical sciences.

A full discussion of random effect interpretations is beyond the present paper's scope. We note, however, that discussions of spatial random effects are generally either unclear about their interpretation or seem to treat them as Scheffé-style random effects. Finally, it is both economical and accurate to describe all non-Scheffé-style random effects as formal devices to implement a smoother, interpreting shrinkage estimation as a kind of smoothing, so from now on we do so.

## 3.1 The random effect S is a formal device to implement a smoother

Consider situations in which $\mathbf{S}$ is not a Scheffé-style random effect. For these situations, we have seen two interpretations of spatial confounding.

- (i) Spatially-correlated errors remove bias in estimating $\beta$ and are generally conservative (Clayton et al 1993).

- (ii) Spatially-correlated errors can introduce or remove bias in estimating $\beta$ and are not necessarily conservative (Wakefield 2007; implicit in Reich et al 2006).

It is commonly argued (e.g., Clayton et al 1993) that introducing spatially correlated errors into a model, as with $\mathbf{S} + \boldsymbol{\epsilon}$, captures the effects of spatially-structured missing covariates and thus adjusts the estimate of $\beta$ for such missing covariates even if we have no idea what those covariates might be. Interpretation (i) reflects this view. We have also heard a somewhat different statement of this view, as in: "I *know* I am missing some confounders, in fact I have some specific confounders in mind that I was unable to collect, but from experience I know they have a spatial pattern. Therefore, I will add $\mathbf{S}$ to the model to try to recover them and let the data decide how much can be recovered." In some fields, it is nearly impossible to get a paper published unless a random effect is included for this purpose.

We can now evaluate this view using the results of Section 2, which are a modest elaboration of linear-model theory. Indeed, the only aspect of the present problem not present in linear-model

11

theory is that most of the canonical coefficients $b_i$ are shrunk toward 0, although the $b_i$ that produce spatial confounding are shrunk the least and thus deviate least from linear-model theory.

To make the discussion concrete, consider estimating $\beta$ using the model $\mathbf{y} = \mathbf{1}_n\alpha + \mathbf{X}\beta + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ is iid normal error, then estimating $\beta$ using a larger model, either $\mathbf{y} = \mathbf{1}_n\alpha + \mathbf{X}\beta + \mathbf{H}\gamma + \boldsymbol{\epsilon}$, where $\mathbf{H}$ is a supposed missing covariate, or using model (2), which adds the ICAR-distributed spatial random effect $\mathbf{S}$. Appendix C gives explicit expressions for the adjustment in the estimate of $\beta$ under either of these larger models, and for the expected adjustment assuming the data were generated by the model

$$\mathbf{y} = \mathbf{1}_n\alpha + \mathbf{X}\beta + \mathbf{H}\gamma + \boldsymbol{\epsilon}. \tag{13}$$

We summarize Appendix C's results here.

There is no necessary relationship between the adjustment to $\beta$'s estimate arising from adding $\mathbf{S}$ to the model and the adjustment arising from adding the supposed missing covariate $\mathbf{H}$. This is most striking if we suppose that $\mathbf{H}$ is uncorrelated with $\mathbf{X}$, so that adding $\mathbf{H}$ to the model would not change the estimate of $\beta$. In this case, adding the spatial random effect $\mathbf{S}$ *does* adjust the estimate of $\beta$, and in a manner that depends not on $\mathbf{H}$ but on the correlation of $\mathbf{X}$ and $\mathbf{y}$ and on the spatial map. If the data are generated by (13), the expected adjustment in $\beta$'s estimate from adding $\mathbf{S}$ to the model is not zero in general and can be biased in either direction. If there *are* no missing covariates, adding $\mathbf{S}$ nonetheless adjusts $\beta$'s estimate in the manner just described although the expected adjustment is zero. It is fair to describe such adjustments as haphazard.

Now suppose $\mathbf{H}$ is correlated with $\mathbf{X}$, so that adding $\mathbf{H}$ to the model changes the estimate of $\beta$. In this case, the adjustment to $\beta$'s estimate under the spatial model is again biased. The bias can be large and either positive or negative, depending on the degree of smoothing (more smoothness generally implies larger bias) and depending haphazardly on $\mathbf{H}$ and on the spatial map. The bias can even be in the wrong direction, so that on average $\beta$'s estimate becomes larger when it would become smaller if $\mathbf{H}$ were added to the model.

Therefore, adding spatially correlated errors is not conservative: a canonical regressor $Z_i$ that is collinear with $\mathbf{X}$ can cause $\beta$'s estimate to increase in absolute value just as in ordinary linear models. Further, in cases in which $\beta$'s estimate should not be adjusted, introducing spatially-correlated errors will, nonetheless, adjust the estimate haphazardly.

From the perspective of linear-model theory, it seems perverse to use an error term to adjust for the possibility of missing confounders. The analog in ordinary linear models would be to move part of the fitted coefficients into error to allow for the possibility of as-yet-unconceived missing confounders. In using an ordinary linear model, we know that if missing confounders are correlated with included fixed effects, variation in $\mathbf{y}$ that would be attributed to the missing confounders is instead attributed to the included fixed effects. We acknowledge that possibility in the standard disclaimer that if we have omitted confounders, our coefficient estimates could be wrong. In spatial modeling, the analogy to this practice would be to use restricted spatial regression, so that all variation in $\mathbf{y}$ in the column space of included fixed effects is attributed to those included effects instead of being haphazardly re-allocated to the spatial random effect.

Therefore, interpretation (i) cannot be sustained and interpretation (ii) is correct, when the random effect **S** is interpreted as a mere formal device to implement a smoother. Adding spatially-correlated errors cannot be expected to capture the effect of a spatially-structured missing covariate, but only to smooth fitted values and discount the sample size in computing standard errors or posterior standard deviations for fixed effects. Therefore, in such cases you should *always* use restricted spatial regression so the sample size can be discounted without distorting the fixed effect estimate. If you are concerned about specific unmeasured confounders, you should add to the model a suitable explicit fixed effect, not adjust haphazardly by means of a spatially-correlated error. Finally, conclusions from such analyses should be qualified as in any other observational study, e.g., we have estimated the association of our outcome with our regressors accounting for measured confounders, and if we have omitted confounders, then our estimate could be wrong.

## 3.2   S is a Scheffé-style random effect

For these situations, we have seen three interpretations.

- (iii) The spatial effect **S** is collinear with the fixed effect, but neither estimate of $\beta$ is biased (David B. Nelson, personal communication).

- (iv) Adding the spatial effect **S** creates information loss, but neither estimate of $\beta$ is biased (David B. Nelson, personal communication).

- (v) Because error is correlated with the regressor **X** in the sense commonly used in econometrics, *both* estimates of $\beta$ are biased (Paciorek 2009).

Interpretations (iii) and (iv) treat the fixed effect **X** as measured without error and not otherwise drawn from a probability distribution ("fixed and known"), while interpretation (v) treats **X** as drawn from a probability distribution. Interpreting spatial confounding therefore depends on whether **X** is interpreted as fixed and known or as a random variable. This is a messy business, which seems to be determined in practice less by facts than by the department in which one was trained. The present authors' training inclines us to view **X** as fixed and known as a default, while econometricians, for example, seem inclined to the opposite default.

To see the difficulty, consider an example in which the random effect is hospitals selected as a random sample from a population of hospitals, and the fixed effect is an indicator of whether a hospital is a teaching hospital. The present authors' default is to treat teaching status as fixed and known. However, if we have drawn 20 hospitals at random, then the teaching status of hospital $i$ is a random variable determined by the sample of hospitals we happen to draw, so **X** is drawn from a probability distribution. But what if, as often happens, sampling is stratified by teaching status to ensure that (say) 10 hospitals are teaching hospitals and 10 are not? Now teaching status is fixed and known. But what if someone gives us the dataset and we don't know whether sampling was stratified by teaching status? One might argue that our ignorance disqualifies us from analyzing these data, but that argument is not compelling to, for example, people who interpret

the Likelihood Principle as meaning they can ignore the sampling mechanism, or to many people who do not have a tenured, hard-money faculty position.

Again, a full discussion of this issue is beyond the present paper's scope. It is also unnecessary for the present purpose, because there are unarguable instances of each kind of fixed effect. An example of a fixed and known $\mathbf{X}$ could arise in analyzing air pollution measured at many fixed monitoring stations on each day in a year. The days could be interpreted as a Scheffé-style random effect, and the elevation of each monitoring station as a fixed and known $\mathbf{X}$. For an example of $\mathbf{X}$ plainly drawn from a probability distribution, consider the hospitals example just above, where the morbidity score for each hospital's patients is a random variable for the period being studied.

So first assume $\mathbf{X}$ is fixed and known. It is then straightforward to show that both (iii) and (iv) are correct and indeed arguably identical, though we think they are worth distinguishing. Interpretation (iv) follows from Section 2.3 and the familiar fact that generalized least squares gives unbiased estimates even when the covariance matrix is specified incorrectly. For interpretation (iii), recall from equation (8) that the estimate of $\beta$ in the spatial model (which, given $\tau_e$ and $\tau_s$, is both the posterior mean and the usual estimate following maximization of the restricted likelihood) is

$$E(\beta|\tau_e, \tau_s, \mathbf{y}) \quad = \hat{\beta}_{OLS} - (n-1)^{-1/2}\mathbf{R}'E(\mathbf{b}|\tau_e, \tau_s, \mathbf{y}). \tag{14}$$

By (7), $(n-1)^{-1/2}\mathbf{R}'E(\mathbf{b}|\tau_e, \tau_s, \mathbf{y})$ can be written as $\mathbf{KP}^c\mathbf{y}$ where $\mathbf{K}$ is an appropriate-sized known square matrix. Recalling that $\mathbf{P}^c = \mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$, $\mathbf{P}^c\mathbf{y} = \mathbf{P}^c(\mathbf{Zb} + \boldsymbol{\epsilon})$, which has expectation 0 with respect to $\mathbf{b}$ and $\boldsymbol{\epsilon}$. Hence the spatial and OLS estimates of $\beta$ have the same expectation and are unbiased based on the aforementioned familiar fact about generalized least squares.

Now assume $\mathbf{X}$ is a random variable. Paciorek (2009) interprets model (5) as $\mathbf{y} = \mathbf{X}\beta + \boldsymbol{\psi}$ where the error term $\boldsymbol{\psi} = \mathbf{Zb} + \boldsymbol{\epsilon}$ has a non-diagonal covariance matrix. Because $\mathbf{X}'\mathbf{Z} \neq 0$, $\mathbf{X}$ is correlated with $\boldsymbol{\psi}$, so by the standard result in econometrics, *both* the OLS estimate of $\beta$ and the estimate of $\beta$ using (5) are biased. Formulating the result this way is more precise than the common statement that bias arises when "the random effect is correlated with the fixed effect", because the distribution of "the random effect" depends on the parameterization: the random effect $\mathbf{b}$ in model (5) is independent of $\mathbf{X}$, but the random effect $\mathbf{S}$ in model (2) is not.

The main point of Paciorek (2009), which presumes the spatial random effect $\mathbf{S}$ captures a missing covariate, is that "bias [in estimating $\beta$] is reduced only when there is variation in the covariate [$\mathbf{X}$] at a scale smaller than the scale of the unmeasured covariate [supposedly captured by $\mathbf{S}$]". We conjecture that this can be interpreted in Section 2.2's terms as meaning that bias is reduced if $\mathbf{X}$ is not too highly correlated with the low-frequency canonical regressors $Z_i$ that have small $d_i$ and hence little shrinkage in $b_i$.

Paciorek (2009) concludes that restricted spatial regression is either irrelevant to the issue of bias in estimating $\beta$ or makes an overly strong assumption by attributing to $\mathbf{X}$ all of the disputed variation in $\mathbf{y}$. The latter appears to presume that the spatial random effect $\mathbf{S}$ captures an unspecified missing covariate, which, we have argued, is difficult to sustain. However, this area of research is just beginning and much remains to be developed.

# 4  Conclusion

The preceding sections laid out the mechanics by which spatial confounding occurs, expanding on Reich et al (2006); showed briefly that this is not an artifact of the ICAR model but is far more general; gave an alternative analysis (restricted spatial regression) that removes spatial confounding; and considered proposed interpretations of spatial confounding, concluding that restricted spatial regression should be used routinely when $\mathbf{S}$ is a formal device to implement a smoother, a common situation.

The literature on spatial confounding is small but it appears many people encounter this problem in practice. Thus although the present paper is not the last word on the subject, it does bring together the various approaches to this phenomenon, which should in time yield generally-accepted advice for statistical practice.

Our understanding of the mechanics of spatial confounding is underdeveloped in certain respects. In debunking the common belief that spatially-correlated errors adjust for unspecified missing covariates, our derivation (Appendix C) took the smoothing ratio $r = \tau_s/\tau_e$ as given. Although $r$'s marginal posterior distribution is easily derived when $\tau_e$ has a gamma prior, it is harder to interpret than the posterior mean of $\beta$ so its implications are as yet unclear. However, it should be possible to extract some generalizations which, with the expressions in Appendix C, will permit understanding of the situations in which spatial confounding will and will not occur. We hypothesize this work will show that whenever both $\mathbf{y}$ and $\mathbf{X}$ are highly correlated with $Z_{n-G}$, the canonical regressor with the least-smoothed coefficient, there will be little smoothing ($r$ will be small) and $\beta$'s estimate under the spatial model will be close to zero. In other words, we hypothesize that in any map, when both $\mathbf{y}$ and $\mathbf{X}$ show a strong trend along the long axis of the map, adding a spatially-correlated error will nullify the obvious association between $\mathbf{y}$ and $\mathbf{X}$ as it did in the Slovenian data.

The theory is particularly underdeveloped for the situation in which both $\mathbf{S}$ and $\mathbf{X}$ can be interpreted as random in Scheffé's sense; Paciorek (2009) is a first step in what should be a rich area of research.

# Appendix A: A small exploration of spatial confounding in some common geostatistical models

In Section 2.5, each municipality $i$ was assigned north-south and east-west coordinates. In the computations below, each coordinate was centered so it averaged zero across the municipalities, and both centered coordinates were divided by a common scaling constant. The symbols $N_i$ and $E_i$ refer to the centered and scaled north-south and east-west coordinates respectively, which had standard deviations 0.76 and 1.24 respectively.

We considered four forms of cov($\mathbf{S}$) determined by combinations of two distance measures and two specifications of correlation between municipalities as a function of distance between them. The measures describing distance between municipalities $i$ and $j$, $\delta_{ij}$, were Euclidean distance $[(N_i - N_j)^2 + (E_i - E_j)^2]^{0.5}$ and maximum distance $\max(|N_i - N_j|, |E_i - E_j|)$, for which the largest

Table 1: Correlation between eigenvectors corresponding to $j^{th}$ smallest eigenvalues of the ICAR's $\mathbf{Q}$ matrix and of $\text{cov}(\mathbf{S})^{-1}$, for four specifications of $\text{cov}(\mathbf{S})$ and various $\theta$.

| distance | function | j | | | | | $\theta$ | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | 0.5 | 1.0 | 1.5 | 2.0 | 2.5 | 3.0 | 3.5 | 4.0 | 4.5 | 5.0 |
| Euclidean | exponential | 2 | 0.89 | 0.95 | 0.95 | 0.96 | 0.96 | 0.96 | 0.96 | 0.96 | 0.96 | 0.96 |
| | | 3 | 0.79 | 0.87 | 0.88 | 0.88 | 0.88 | 0.87 | 0.87 | 0.86 | 0.86 | 0.85 |
| | | 4 | 0.42 | 0.20 | 0.21 | 0.23 | 0.24 | 0.25 | 0.25 | 0.26 | 0.26 | 0.26 |
| Euclidean | linear | 2 | 0.34 | 0.61 | 0.88 | 0.92 | 0.94 | 0.95 | 0.96 | 0.96 | 0.96 | 0.96 |
| | | 3 | 0.44 | 0.60 | 0.80 | 0.87 | 0.88 | 0.85 | 0.66 | 0.28 | 0.18 | 0.19 |
| | | 4 | 0.25 | 0.33 | 0.10 | 0.15 | 0.21 | 0.27 | 0.28 | 0.20 | 0.15 | 0.13 |
| maximum | exponential | 2 | 0.90 | 0.94 | 0.95 | 0.95 | 0.95 | 0.95 | 0.95 | 0.95 | 0.95 | 0.95 |
| | | 3 | 0.81 | 0.86 | 0.87 | 0.87 | 0.87 | 0.86 | 0.86 | 0.85 | 0.84 | 0.83 |
| | | 4 | 0.17 | 0.11 | 0.10 | 0.10 | 0.09 | 0.09 | 0.09 | 0.08 | 0.08 | 0.08 |
| maximum | linear | 2 | 0.36 | 0.72 | 0.90 | 0.93 | 0.94 | 0.94 | 0.95 | 0.95 | 0.95 | 0.95 |
| | | 3 | 0.15 | 0.66 | 0.83 | 0.87 | 0.87 | 0.87 | 0.26 | 0.25 | 0.26 | 0.29 |
| | | 4 | 0.26 | 0.36 | 0.09 | 0.11 | 0.14 | 0.13 | 0.01 | 0.03 | 0.04 | 0.04 |

distances between two municipalities were 5.4 and 4.7 respectively. The forms specifying correlation as a function of distance were exponential, $\exp(-\delta_{ij}/\theta)$, and linear, $\max[(1 - \delta_{ij}/\theta), 0]$.

For each form of $\text{cov}(\mathbf{S})$, Table 1 shows the correlation between the eigenvector corresponding to the $j^{th}$ smallest eigenvalue of $\text{cov}(\mathbf{S})^{-1}$ and the eigenvector corresponding to the $j^{th}$ smallest eigenvalue of the ICAR model's neighbor matrix $\mathbf{Q}$, for a range of values of $\text{cov}(\mathbf{S})$'s tuning constant $\theta$. When this correlation is high for $j = 2$, the geostatistical specification for $\text{cov}(\mathbf{S})$ will produce the same spatial confounding as the ICAR model.

Table 1 shows that generally the correlations between eigenvectors of $\text{cov}(\mathbf{S})^{-1}$ and $\mathbf{Q}$ are high for $j = 2$ and $j = 3$, but fall off substantially for $j = 4$ and for $j > 4$ (not shown). Results for the four $\text{cov}(\mathbf{S})$ differ somewhat in details. The two distance measures behave similarly. However, while the exponential function of distance gives high correlations for all values of $\theta$ shown here for $j = 2, 3$, the linear function of distance shows smaller correlations for small values of $\theta$, most likely because the spatial correlation dies off so quickly for small $\theta$, and for $j = 3$ for large $\theta$ as well.

## Appendix B: Spatial confounding with spectral methods

Section 2.2's results for a discrete spatial domain can be extended to Gaussian process models defined on a continuous spatial domain. Let $y(\text{s}_i) = \text{x}_i'\boldsymbol{\beta} + S(\text{s}_i) + \varepsilon_i$, where $\text{s}_i \in \mathcal{R}^2$, $\varepsilon_i \overset{iid}{\sim} \text{N}(0, \tau_e)$ is pure error, and $S$ is a spatial process with mean zero, precision $\tau_s$, and stationary spatial correlation function $\text{Cor}(\text{s}_i, \text{s}_j) = \rho(\text{s}_i - \text{s}_j)$.

16

The spectral representation theorem states that $S(\mathrm{s})$ can be written

$$S(\mathrm{s}) = \int_{\mathcal{R}^2} \cos(\boldsymbol{\omega}'\mathrm{s})db_1(\boldsymbol{\omega}) + \int_{\mathcal{R}^2} \sin(\boldsymbol{\omega}'\mathrm{s})db_2(\boldsymbol{\omega}), \tag{15}$$

where $\boldsymbol{\omega} = (\omega_1, \omega_2)' \in \mathcal{R}^2$ is a frequency and $b_1$ and $b_2$ are independent Gaussian processes with mean zero, orthogonal increments, and $E(|db_j(\boldsymbol{\omega})|^2) = F(\boldsymbol{\omega})/\tau_s$. The spectral representation formulates the spatial process as a convolution of trigonometric basis functions and stochastic processes in the frequency domain with independent increments. The spatial correlation $\rho$ is directly related to the spectral density $F$:

$$\rho(\mathrm{s}_i - \mathrm{s}_j) = \int_{\mathcal{R}^2} \cos[\boldsymbol{\omega}'(\mathrm{s}_i - \mathrm{s}_j)]dF(\boldsymbol{\omega}). \tag{16}$$

The spectral density is often decreasing in $||\boldsymbol{\omega}||$, for example, $F(\boldsymbol{\omega}) \propto \exp(-||\boldsymbol{\omega}||^2/(4\phi))$ corresponds to the squared-exponential covariance $\rho(\mathrm{s}_i - \mathrm{s}_j) = \exp(-\phi||\mathrm{s}_i - \mathrm{s}_j||^2)$.

Assuming the observations lie on a $m \times m$ square grid with distance one between neighboring sites, $S$'s density can be approximated (Whittle 1954) by a sum over a finite grid of Fourier frequencies $\boldsymbol{\omega}_l \in \{2\pi\lfloor-(m-1)/2\rfloor/m, ..., 2\pi(n - \lfloor n/2\rfloor)/m\}^2$,

$$S(\mathrm{s}) = \sum_{j=1}^{2}\sum_{l=1}^{m^2} Z_j(\mathrm{s}, \boldsymbol{\omega}_l)b_{jl} \tag{17}$$

$$b_{jl} \sim \mathrm{N}(0, \tau_s d_l)$$

where $\lfloor x \rfloor$ is the smallest integer greater than or equal to $x$, $Z_1(\mathrm{s}, \boldsymbol{\omega}) = \cos(\boldsymbol{\omega}'\mathrm{s})$, $Z_2(\mathrm{s}, \boldsymbol{\omega}) = \sin(\boldsymbol{\omega}'\mathrm{s})$, and $\tau_s d_l$ is a precision with $d_l = 1/F(\boldsymbol{\omega}_l)$. In this special case, the $Z_j$ are orthogonal, i.e., $\sum_{l=1}^{m^2} Z_j(\mathrm{s}, \boldsymbol{\omega}_l)Z_k(\mathrm{s}', \boldsymbol{\omega}_l) = I(k = j)I(\mathrm{s} = \mathrm{s}')$. This approximation may induce edge and aliasing effects in the spatial covariance, especially for small grids, but the approximation is useful for studying the fixed effects. This representation is directly analogous to the ICAR model as in (5), so the rest of the analysis applied to the ICAR model follows directly.

Comparing the spectral representation (17) with the ICAR model (5), the trigonometric functions $Z_1$ and $Z_2$ are analogous to Section 2.2's eigenvectors $Z_j$, and the inverse of the spectral density $1/F(\omega)$ is analogous to the eigenvalues $d_j$ in Section 2.2. Unlike the eigenvectors and eigenvalues of the ICAR model, $Z_1$, $Z_2$ and $F$ have explicit forms, so the role of the covariate's spatial scale is more clear. High-frequency terms (large $||\boldsymbol{\omega}||$) have small prior variance (small $F(||\boldsymbol{\omega}||)$) and are shrunk a great deal, while low-frequency terms (small $||\boldsymbol{\omega}||$) have large prior variance (large $F(||\boldsymbol{\omega}||)$) and are shrunk relatively little. In this parameterization, the correlations between a covariate and the $Z_j$ clearly describe the covariate's variation at different spatial scales.

# Appendix C: How $\beta$'s estimate changes when a supposed missing covariate or the random effect S are added to the model

Consider these three models:

- **Model 0**: $\mathbf{y} = \mathbf{1}_n \alpha + \mathbf{X}\beta + \boldsymbol{\epsilon}$

- **Model H**: $\mathbf{y} = \mathbf{1}_n \alpha + \mathbf{X}\beta + \mathbf{H}\gamma + \boldsymbol{\epsilon}$

- **Model S**: $\mathbf{y} = \mathbf{X}\beta + \mathbf{Zb} + \boldsymbol{\epsilon}$,

where Model S, the spatial model, is the same as (2) and (5). Assume, as before, that $\mathbf{X}$ is centered and scaled so $\mathbf{1}'_n\mathbf{X} = 0$ and $\mathbf{X}'\mathbf{X} = n-1$; assume $\mathbf{H}$ is centered and scaled the same way; and assume $\mathbf{y}$ is centered but not scaled. Bayesian results below assume flat (improper) priors on $\alpha$, $\beta$, and $\gamma$.

Under Model 0, the estimate of $\beta$ — the posterior mean or least-squares estimate — is

$$\hat{\beta}^{(0)} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'y = \rho_{XY}(\mathbf{y}'\mathbf{y})^{0.5} \tag{18}$$

where $\rho_{AB}$ is Pearson's correlation of the vectors $\mathbf{A}$ and $\mathbf{B}$. Under Models H and S, the estimates of $\beta$ given $\tau_s$ and $\tau_e$ — the conditional posterior mean for both models; for Model H this is also the least-squares estimate, while for Model S it is the customary estimate after maximizing the restricted likelihood — are:

$$\hat{\beta}^{(H)} = \hat{\beta}^{(0)} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{H}(\mathbf{H}'\mathbf{P}^c\mathbf{H})^{-1}\mathbf{H}'\mathbf{P}^c\mathbf{y} \tag{19}$$

$$\hat{\beta}^{(S)} = \hat{\beta}^{(0)} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Z}(\mathbf{Z}'\mathbf{P}^c\mathbf{Z} + r\mathbf{D})^{-1}\mathbf{Z}'\mathbf{P}^c\mathbf{y}, \tag{20}$$

where, as before, $r = \tau_s/\tau_e$, $\mathbf{P}^c = \mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$, and $\mathbf{D} = diag(d_1, \ldots, d_{n-G}, 0, \ldots, 0)$ is the diagonal matrix containing the eigenvalues $d_i$ of the spatial-neighbor matrix $\mathbf{Q}$. The estimate $\hat{\beta}^{(S)}$ was given in Section 2.2 and $\hat{\beta}^{(H)}$ is derived by a similar argument.

Define $\mathbf{B}_X = \mathbf{Z}'\mathbf{X}/(n-1)^{0.5}$; the entries in $\mathbf{B}_X$ are the correlations between $\mathbf{X}$ and the columns of $\mathbf{Z}$ (we called this $\mathbf{R}$ in Section 2.2). Define $\mathbf{B}_H$ analogously as $\mathbf{B}_H = \mathbf{Z}'\mathbf{H}/(n-1)^{0.5}$. Finally, define $\mathbf{B}_y = \mathbf{Z}'\mathbf{y}/[(n-1)(\mathbf{y}'\mathbf{y})]^{0.5}$; $\mathbf{B}_y$'s entries are the correlations between $\mathbf{y}$ and the columns of $\mathbf{Z}$. Then the estimates of $\beta$ under Models H and S, given $\tau_s$ and $\tau_e$, can be shown to be

$$\hat{\beta}^{(H)} = \hat{\beta}^{(0)}\left[1 - \frac{\frac{\rho_{XH}\rho_{HY}}{\rho_{XY}} - \rho_{XH}^2}{1 - \rho_{XH}^2}\right] \tag{21}$$

$$\hat{\beta}^{(S)} = \hat{\beta}^{(0)}\left[1 - \frac{\frac{\rho'_{XY}}{\rho_{XY}} - q}{1 - q}\right],$$

where $\rho'_{XY} = \mathbf{B}'_X(\mathbf{I} + r\mathbf{D})^{-1}\mathbf{B}_y = \mathbf{X}'(I + r\mathbf{Q})^{-1}\mathbf{y}/((n-1)\mathbf{y}'\mathbf{y})^{0.5}$ and $q = \mathbf{B}'_X(\mathbf{I} + r\mathbf{D})^{-1}\mathbf{B}_X = \mathbf{X}'(I + r\mathbf{Q})^{-1}\mathbf{X}/(n-1)$.

In (21), the expressions in square brackets for $\hat{\beta}^{(H)}$ and for $\hat{\beta}^{(S)}$ have no necessary relation to each other. For example, if $\mathbf{X}$ and $\mathbf{H}$ are uncorrelated, so $\rho_{XH} = 0$, then $\hat{\beta}^{(H)} = \hat{\beta}^{(0)}$ but $\hat{\beta}^{(S)} \neq \hat{\beta}^{(0)}$. In particular, $\hat{\beta}^{(S)}$ can be larger or smaller than $\hat{\beta}^{(H)}$ in absolute value; this depends on how $(\mathbf{I}+r\mathbf{D})^{-1}$ differentially downweights specific coordinates of $\mathbf{B}_X$ and $\mathbf{B}_y$. When $\rho'_{XY} \approx \rho_{XY}$, $\hat{\beta}^{(S)} \approx 0$. This happens if the $i^{th}$ coordinate of $\mathbf{B}_X$, the correlation of $\mathbf{X}$ and $Z_i$, is large; the $i^{th}$ coordinate of $\mathbf{B}_y$, the correlation of $\mathbf{y}$ and $Z_i$, is large; and $d_i$ and $r$ are small, as in the Slovenian data.

If the data are generated by Model H, we can treat $\hat{\beta}^{(H)}$ and $\hat{\beta}^{(S)}$ as functions of $\mathbf{y}$, holding the precisions $\tau_e$ and $\tau_s$ fixed, and compute the expected change in $\beta$'s estimate from adding either $\mathbf{H}$ or $\mathbf{S}$ to Model 0. The expected changes are:

$$E(\hat{\beta}^{(H)} - \hat{\beta}^{(0)}|\tau_e) = \rho_{XH}\gamma \tag{22}$$

$$E(\hat{\beta}^{(S)} - \hat{\beta}^{(0)}|\tau_e, \tau_s) = \left[ \frac{\frac{\rho'_{XH}}{\rho_{XH}} - q}{1 - q} \right] \rho_{XH}\gamma \quad \text{if } \rho_{XH} \neq 0 \tag{23}$$

$$= \frac{\rho'_{XH}}{1 - q}\gamma \quad \text{if } \rho_{XH} = 0, \tag{24}$$

where $\rho'_{XH} = \mathbf{B}'_X(\mathbf{I} + r\mathbf{D})^{-1}\mathbf{B}_H = \mathbf{X}'(I + r\mathbf{Q})^{-1}\mathbf{H}/(n-1)$. If $\gamma \neq 0$, the expected adjustment under Model S is biased, with the bias depending on how $(\mathbf{I} + r\mathbf{D})^{-1}$ differentially downweights specific coordinates of $\mathbf{B}_X$ and $\mathbf{B}_H$; the bias can be positive or negative. If $\mathbf{H} = Z_j$, then $\rho'_{XH} = \rho_{XH}/(1 + rd_j)$, so (23) becomes

$$E(\hat{\beta}^{(S)} - \hat{\beta}^{(0)}|\tau_e, \tau_s) = \left[ \frac{(1 + rd_j)^{-1} - q}{1 - q} \right] \rho_{XH}\gamma. \tag{25}$$

The expression in square brackets is less than 1 when $rd_j > 0$ and becomes negative if $r$ or $d_j$ is large enough, i.e., Model S adjusts $\hat{\beta}$ in the wrong direction from $\hat{\beta}^{(0)}$. If there is no missing covariate, $\gamma = 0$, then the expected adjustment is zero under both Model H and Model S.

## Acknowledgements

# References

Clayton DG, Bernardinelli L, Montomoli C (1993). Spatial correlation in ecological analysis. *Int. J. Epi.*, **22**,1193-1201.

Hodges JS, Carlin BP, Fan Q (2003). On the precision of the conditionally autoregressive prior in spatial models. *Biometrics*, **59**, 317-322.

Paciorek CJ (2009). The importance of scale for spatial-confounding bias and precision of spatial regression estimators. Manuscript dated March 12, 2009.

Reich BJ, Hodges JS (2008). Identification of the variance components in the general two-variance linear model. *J. Stat. Plan. Inf.*, **138**, 1592-1604.

Reich BJ, Hodges JS, Zadnik V (2006). Effects of residual smoothing on the posterior of the fixed effects in disease-mapping models. *Biometrics*, **62**, 1197-1206. Errata available from the corresponding author.

Ruppert D, Wand MP, Carroll RJ (2003). *Semiparametric Regression.* New York:Cambridge.

Salkowski NJ (2008). Using the SemiPar Package. Class project, available at URL http://www.biostat.umn.edu/~hodges/SalkowskiRPMProject.pdf.

Scheffé H (1959). *The Analysis of Variance*, New York:Wiley.

Spiegelhalter DM, Best DG, Carlin BP, Van der Linde A (2002). Bayesian measures of model complexity and fit (with discussion). *J. Royal Statistical Society, Ser. B*, **64**, 583-639.

Steinberg DM, Bursztyn D (2004). Data analytic tools for understanding random field regression models. *Technometrics*, **46**, 411-420.

Wakefield J (2007). Disease mapping and spatial regression with count data. *Biostatistics*, **8**, 158-183.

Whittle P (1954). On stationary processes in the plane. *Biometrika*, **41**, 434-449.

Zadnik V, Reich BJ (2006). Analysis of the relationship between socioeconomic factors and stomach cancer incidence in Slovenia. *Neoplasma*, **53**, 103-110.

Figure 1: For the Slovenian municipalities, panel (a): Observed standardized incidence ratio $SIR = O_i/E_i$; panel (b): Centered and scaled socioeconomic status $SEc$.

(a) $SIR$
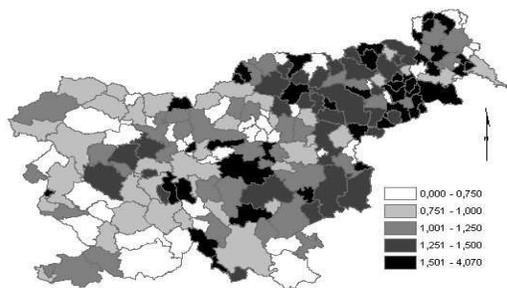
(b) $SEc$
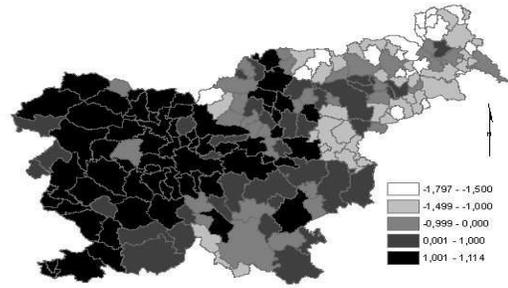
Figure 2: Two canonical regressors, columns of the matrix $\mathbf{Z}$; panel (a): $Z_{193}$, with the smallest positive eigenvalue $d_{193}$ and thus the least-shrunk coefficient $b_{193}$ among the shrunk coefficients; panel (b): $Z_1$, with the largest eigenvalue $d_{193}$ and thus the most-shrunk coefficient $b_1$.

(a) $Z_{193}$                                    (b) $Z_1$