



JOHNS HOPKINS
BLOOMBERG
SCHOOL of PUBLIC HEALTH

Johns Hopkins University, Dept. of Biostatistics Working Papers

12-1-2008

Spatial Misalignment in time series studies of air pollution and health data

Roger D. Peng

Johns Hopkins University, rpeng@jhsph.edu

Michelle L. Bell

Suggested Citation

Peng, Roger D. and Bell, Michelle L., "Spatial Misalignment in time series studies of air pollution and health data" (December 2008). *Johns Hopkins University, Dept. of Biostatistics Working Papers*. Working Paper 176. <http://biostats.bepress.com/jhubiostat/paper176>

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

Copyright © 2011 by the authors

Spatial misalignment in time series studies of air pollution and health data

Roger D. Peng Michelle L. Bell

December 1, 2008

Abstract

Time series studies of environmental exposures often involve comparing daily changes in a toxicant measured at a point in space with daily changes in an aggregate measure of health. Spatial misalignment of the exposure and response variables can bias the estimation of health risk and the magnitude of this bias depends on the spatial variation of the exposure of interest. In air pollution epidemiology, there is an increasing focus on estimating the health effects of the chemical components of particulate matter. One issue that is raised by this new focus is the spatial misalignment error introduced by the lack of spatial homogeneity in many of the particulate matter components. Current approaches to estimating short-term health risks via time series modeling do not take into account the spatial properties of the chemical components and therefore could result in biased estimation of those risks. We present a spatial-temporal statistical model for quantifying spatial misalignment error and show how adjusted health risk estimates can be obtained using a regression calibration approach and a two-stage Bayesian model. We apply our methods to a database containing information on hospital admissions, air pollution, and weather for 20 large urban counties in the United States.

1 Introduction

Estimating the health risks of environmental exposures often involves examining data at different levels of spatial resolution. This mismatch between data measured at different resolutions results

in spatial misalignment (Banerjee et al., 2004) which can induce error and bias estimates of risk. Spatial misalignment in environmental health studies is common because the exposure data and the health data often come from independent sources. For example, in an air pollution study in the US, data on ambient air pollution levels often is based on a network of monitors operated by the US Environmental Protection Agency (EPA) where each monitor measures pollution at a specific point location. Data on an outcome of interest, such as the numbers of hospital admissions for cardiovascular disease, might come from the Centers for Medicare and Medicaid Services. In cohort studies which compare health outcomes and exposures across people, information about individual persons may be available but exposure data may be derived from a computer model at a much lower resolution. Because health and exposure data are often collected independently of each other, they are rarely spatially aligned. Hence, a direct comparison of the exposure and health outcome is not possible without a model (or an assumption) to align the two in the spatial domain.

1.1 Air pollution and health

In a time series study of air pollution and health, one is interested in estimating associations between daily changes in county-wide hospital admissions or mortality counts with daily changes in county-wide average levels of a specific pollutant. The problem is that we do not directly observe county-wide average pollutant levels. Rather, we have measurements taken at a handful of monitors (sometimes only one) located inside the county boundaries. For a spatially homogeneous pollutant, the value of the pollutant at a single monitor can be representative of the county-wide average ambient level of that pollutant. Some pollutants, particularly some gases such as ozone, are reasonably spatially homogeneous across the area of a county. The total mass of particulate matter (PM) less than $2.5 \mu\text{m}$ in aerodynamic diameter ($\text{PM}_{2.5}$), whose health risks have been examined extensively, is fairly spatially homogeneous and monitor measurements of $\text{PM}_{2.5}$ in counties with multiple monitors tend to be highly correlated across both time and space (Peng et al., 2008; Bell et al., 2006). With a pollutant such as $\text{PM}_{2.5}$ the misalignment between the continuous nature of the pollutant process and the aggregated nature of the health data does not typically pose as serious a problem as some other pollutants. In this situation, current approaches for data analysis are generally adequate.

Recently, data have become available from the US EPA's Chemical Speciation Trends Network (STN) as well as State and Local Air Monitoring Stations which provide daily mass concentrations of approximately 60 different chemical elements of $PM_{2.5}$. These data are monitored in over 200 locations around the US starting from the year 2000. Although the data are promising and will be a source of substantial new research, they also raise new statistical challenges. In particular, the usual assumption of spatial homogeneity of the exposure pollutant is less tenable when examining certain components of $PM_{2.5}$ (Bell et al., 2006).

Much of the spatial heterogeneity in the chemical components measured by the STN can be explained by the nature of the sources of the various components. For example, elemental carbon (EC) and organic carbon matter (OCM) tend to be emitted primarily from vehicle or other mobile sources and thus their spatial distribution can depend on the localized nature of those sources. Secondary pollutants such as sulfate and nitrate are created in the air by the chemical and physical transformation of other pollutants and tend to be more regional in nature. Hence, for spatially heterogeneous pollutants such as EC or OCM, the daily level of those pollutants at a single monitor may be a poor surrogate for the daily county-wide average ambient level of that pollutant.

Although we focus mainly on time series studies in this paper, spatial misalignment can also induce error in spatial (cross-sectional) studies of air pollution and health. Gryparis et al. (2008) demonstrate in detail how to handle the misalignment errors in these types of studies and compare the performance of a number of different statistical approaches. An alternate modeling approach has been proposed by Fuentes et al. (2006) for estimating the spatial association between speciated fine particles and mortality. Both approaches introduce a spatial model for the monitored pollutant concentrations and either predict pollutant values at unobserved locations or compute area averages over counties to link with county-level health data. While much work has been done illustrating the problem of spatial misalignment for cross-sectional studies, little has been done for time series studies estimating short-term health effects.

With the emergence of the PM components data and the critical interest in estimating the health effects of those components, the problem of spatial misalignment becomes more important and new statistical methods need to be developed to properly analyze these data. Specifically, there is a

need to develop approaches that account for the varying amounts of spatial information we have about chemical component levels over wide regions. These methods could be used to estimate and report health risks associated with PM components while simultaneously incorporating the spatial information (or lack thereof) we have for each component.

In this paper we describe a general method for estimating health risks associated with PM components from time series models while adjusting for potential spatial misalignment error. We first develop a spatial-temporal model for the exposure of interest and estimate the degree of spatial misalignment error for each component in a location. We then apply two methods—a regression calibration procedure and a two-stage Bayesian model—to estimate the health risks associated with these components and compare the results to standard approaches. Our methods are applied to a large database containing information on daily hospital admissions for cardiovascular diseases and chemical components of ambient particulate matter for 20 large urban counties in the United States.

2 Current Methods

Current approaches to time series analysis of air pollution and health data typically use county-level health data and treat the day as the temporal unit of analysis. On a given day, the outcome is the count of the number of hospitalizations for a specific disease occurring in that county or perhaps the number of deaths due to a specific cause. The exposure is typically the level of a pollutant recorded at a monitor located within the county boundaries on that day. If there are multiple monitors located in the county (sometimes between 2 and 10), then an average of the available monitors on that day is used as a single exposure concentration. For example, in the National Mortality, Morbidity, and Air Pollution Study (NMMAPS), multiple monitors were averaged using a 10% trimmed mean to remove any outlying large or small values (Samet et al., 2000a). Similar approaches have been taken in other time series studies (Katsouyanni et al., 2001).

For a given county there is a time series for the outcome y_t and a single time series for the exposure pollutant \bar{w}_t which might represent the daily averages across multiple monitors. Generalized additive models are often used to estimate the day-to-day association between the two time series while controlling for the potential confounding effects of weather, season, and other factors (Peng et al.,

2006; Welty and Zeger, 2005; Touloumi et al., 2004). Log-linear Poisson models are commonly used with smooth functions of time, temperature, and dew point temperature (or relative humidity),

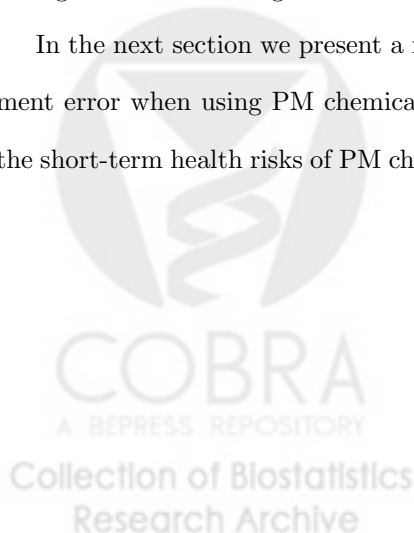
$$\log \mathbb{E}[Y_t] = \alpha + \theta \bar{w}_t + \boldsymbol{\beta}' \mathbf{z}_t + s(t, \lambda_1) + s(\text{temp}_t, \lambda_2) + s(\text{dewpt}_t, \lambda_3) \quad (1)$$

where Y_t is the count of the number of events (i.e. admissions, deaths), \mathbf{z}_t represents a vector of covariates (e.g. indicators for the day of the week, age category-specific intercepts), s is a smooth function, and the λ_i s are smoothing parameters controlling the smoothness of their respective functions. Common choices for the smooth functions include penalized splines, smoothing splines, and parametric natural splines. Variations on the model in (1) are used, depending on meteorological or other local conditions.

The parameter of interest is θ , the log-relative risk of the exposure \bar{w}_t , with the remaining elements of the model being nuisance parameters. When total mass particulate matter is the exposure of interest the risk is often reported as $100 \times (\exp(10\theta) - 1)$, which is the percent increase in the outcome for a $10 \mu\text{g}/\text{m}^3$ increase in PM. Another common increment for reporting risk estimates is the interquartile range (IQR) of the pollutant data.

One key assumption made by current approaches is that the monitor value (or average of monitor values) represents the average ambient concentration. For some pollutants such as $\text{PM}_{2.5}$ total mass this assumption is approximately true. However, for many components of $\text{PM}_{2.5}$, the observed spatial heterogeneity of the component raise concerns about whether the monitor values are good surrogates for the average ambient concentration.

In the next section we present a model for assessing the degree of spatial misalignment measurement error when using PM chemical components data and an approach for adjusting estimates of the short-term health risks of PM chemical components obtained from time series regression models.



3 Statistical Model for Spatial Misalignment

3.1 Spatial–temporal model for exposure

We assume that for a given time t and point location \mathbf{s} , a pollutant can be modeled by a spatial stochastic process

$$w(\mathbf{s}, t) = \mu(\mathbf{s}, t) + \varepsilon(\mathbf{s}, t)$$

where $\mu(\mathbf{s}, t)$ is a fixed effect term and $\varepsilon(\mathbf{s}, t)$ is a mean zero Gaussian process with variance σ^2 and correlation function $\rho(\cdot, \cdot)$. We further assume an isotropic covariance model so that

$$\text{Cov}(\varepsilon(\mathbf{s}, t), \varepsilon(\mathbf{s}', t')) = \begin{cases} \sigma^2 \rho(\|\mathbf{s} - \mathbf{s}'\|) & t = t' \\ 0 & t \neq t' \end{cases}$$

where $\|\cdot\|$ is the Euclidean distance between two points. In this formulation we do not model the temporal correlation structure and assume the data are independent in time conditional on the fixed effect $\mu(\mathbf{s}, t)$, which may contain nonparametric smoothers of time and/or space. Previous studies have shown that after adjustment for seasonal variation and temperature, there is relatively little autocorrelation left in the data. Furthermore, given that the data in which we are most interested are sampled only once every six days, we expect the residual autocorrelation to be small. For the correlation function ρ , we use the flexible Matérn correlation function with parameters ϕ and κ , which has the form

$$\rho(u; \phi, \kappa) = \frac{1}{2^{\kappa-1} \Gamma(\kappa)} \left(\frac{u}{\phi}\right)^{\kappa} \mathcal{K}_{\kappa} \left(\frac{u}{\phi}\right)$$

for $\phi > 0$, $\kappa > 0$, where \mathcal{K} is the modified Bessel function of the third kind. We make use of the `geoR` R package implementation of this model (Ribeiro and Diggle, 2001).

Let $\mathbf{w}_t = (w(\mathbf{s}_1, t), \dots, w(\mathbf{s}_n, t))$ and $\boldsymbol{\mu}_t = (\mu(\mathbf{s}_1, t), \dots, \mu(\mathbf{s}_n, t))$ where $\mathbf{s}_1, \dots, \mathbf{s}_n$ are the locations of all the monitors in the pollutant monitoring network. Then for a given time point t , the observed data at that time point follow the distribution $\mathbf{w}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \sigma^2 H(\phi, \kappa))$, where H is an $n \times n$ correlation matrix with elements

$$[H(\phi, \kappa)]_{ij} = \rho(\|\mathbf{s}_i - \mathbf{s}_j\|; \phi, \kappa). \tag{2}$$

and ρ is the Matérn correlation function. The joint likelihood for the data across all time points is

then

$$L(\sigma, \phi, \kappa) \propto \prod_{t=1}^T \sigma^{-1} |H(\phi, \kappa)|^{-1/2} \exp\left(-\frac{1}{2\sigma^2} (\mathbf{w}_t - \boldsymbol{\mu}_t)' H(\phi, \kappa)^{-1} (\mathbf{w}_t - \boldsymbol{\mu}_t)\right) \quad (3)$$

where the matrix $H(\phi, \kappa)$ is defined as in (2). The likelihood in (3) can be maximized using standard nonlinear optimization techniques to obtain the maximum likelihood estimates for the parameters σ , ϕ , and κ as well as any parameters incorporated into $\boldsymbol{\mu}_t$. To obtain standard error estimates we use the diagonal of the inverse Hessian matrix calculated at the maximum (Nocedal and Wright, 1999).

3.2 Spatial misalignment error model

On a given day t , a monitor in a county can be thought of as providing a surrogate measurement for the county-wide average ambient concentration (for now assume there is only one monitor located in the county). We call this observed monitor value \bar{w}_t , which is the concentration of the component on day t at the monitor location. Let x_t be the true but unobserved county-wide average concentration of the pollutant on day t . The difference between the two values can be described with a classical measurement error model, so that

$$\bar{w}_t = x_t + u_t \quad (4)$$

where u_t is a random variable with $\mathbb{E}[u_t] = \eta_t$ and $\text{Var}(u_t) = \tau_u^2$. The extent to which \bar{w}_t differs from x_t is the spatial misalignment error and τ_u^2 is the error variance. In practice, we occasionally have more than one monitor in a county, so that \bar{w}_t reflects the average of the available monitors on day t . As the number of monitors increases, we would expect τ_u^2 to decrease. Note that the nonzero mean η_t for u_t arises from the inclusion of fixed effects $\mu(\mathbf{s}, t)$ in the spatial-temporal model for the pollutant $w(\mathbf{s}, t)$. If desired, this could be removed by first detrending all of the data but it plays no role in our analysis.

If a pollutant is inherently spatially homogeneous, then τ_u^2 will likely be small and \bar{w}_t will generally serve as a good surrogate for the true county-wide average x_t , even with just a single monitor. If the pollutant is inherently spatially heterogeneous then τ_u^2 will likely be large and \bar{w}_t serves as a poor surrogate for the county-wide average. PM components such as elemental carbon and silicon, as well as the coarse fraction of PM, tend to fall into this latter category.

The classical measurement error model appears appropriate for this situation because we would expect pollutant values at an individual monitor to be more variable over time than the county-wide average. The other assumption made by the classical model is that the errors are independent of x_t . We will examine this assumption further in the data analysis. Ultimately, it may be that neither the classical nor the Berkson model truly describes the relationship between the observed data and the underlying county-wide average, but the classical model seems reasonable in this application.

Given that a single monitor value can be a poor surrogate for the county-wide average ambient concentration, our goal is to use information from other monitors in neighboring counties to obtain a better estimate of the county-wide average. With a spatial-temporal model for the underlying pollutant process, we can estimate the county-wide average concentration on each day and subsequently use those estimated concentrations to obtain health risk estimates for the pollutant.

3.3 Estimating misalignment error

Suppose we have a county represented by polygon A with monitors located at coordinates $\mathbf{v}_1, \dots, \mathbf{v}_m$ within the county boundaries. Then

$$\begin{aligned}\bar{w}_t &= \frac{1}{m} \sum_{i=1}^m w(\mathbf{v}_i, t) \\ x_t &= \frac{1}{\|A\|} \int_A w(\mathbf{s}, t) d\mathbf{s}\end{aligned}$$

Note that the locations $\mathbf{v}_1, \dots, \mathbf{v}_m$ will be a subset of all the locations $\mathbf{s}_1, \dots, \mathbf{s}_n$ used to fit the spatial model for $w(\mathbf{s}, t)$ above.

In a typical air pollution application the number of monitors m in a county might range from 1 to 10. One intermediate target of inference is the misalignment error variance $\tau_u^2 = \text{Var}(u_t) = \text{Var}(\bar{w}_t - x_t)$ for a given county. The monitor values inside the county $w(\mathbf{v}_1, t), \dots, w(\mathbf{v}_m, t)$ and the true county-wide average x_t have a joint Normal distribution

$$\begin{pmatrix} w(\mathbf{v}_1, t) \\ \vdots \\ w(\mathbf{v}_m, t) \\ x_t \end{pmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \boldsymbol{\mu}_t \\ \mu_{x,t} \end{bmatrix}, \sigma^2 \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} \right) \quad (5)$$

where

$$\mu_{x,t} = \frac{1}{\|A\|} \int_A \mu(\mathbf{s}, t) ds$$

and

$$\begin{aligned} [H_{11}]_{ij} &= \rho(\|\mathbf{v}_i - \mathbf{v}_j\|; \phi, \kappa) \\ [H_{12}]_i &= \frac{1}{\|A\|} \int \rho(\|\mathbf{v}_i - \mathbf{s}\|; \phi, \kappa) ds \\ H_{21} &= H'_{12} \\ H_{22} &= \frac{1}{\|A\|^2} \iint \rho(\|\mathbf{s} - \mathbf{s}'\|; \phi, \kappa) ds ds' \end{aligned}$$

The spatial misalignment error variance τ_u^2 for the county can then be calculated as

$$\begin{aligned} \tau_u^2 &= \sigma^2 \left[\frac{1}{m^2} \sum_{i,j} [H_{11}]_{ij} - \frac{2}{m} \sum_i [H_{12}]_i + H_{22} \right] \\ &= \sigma^2 \left[\frac{1}{m^2} \sum_{i,j} \rho(\|\mathbf{v}_i - \mathbf{v}_j\|; \phi, \kappa) - \frac{2}{m} \sum_i \frac{1}{\|A\|} \int \rho(\|\mathbf{v}_i - \mathbf{s}\|; \phi, \kappa) ds \right. \\ &\quad \left. + \frac{1}{\|A\|^2} \iint \rho(\|\mathbf{s} - \mathbf{s}'\|; \phi, \kappa) ds ds' \right] \end{aligned}$$

The values of σ^2 , ϕ , and κ are unknown and so we plug in the maximum likelihood estimates of those parameters to obtain our estimate of τ_u^2 . Because the integrals required above are all over the domain of the county boundary, which is likely to be highly irregular, we use a Monte Carlo approximation. We generate random variables p_1, p_2, \dots, p_B which have a uniform distribution over the area A and then calculate

$$\begin{aligned} \frac{1}{\|A\|} \int \rho(\|\mathbf{s}_i - \mathbf{s}\|; \phi, \kappa) ds &\approx \frac{1}{B} \sum_{j=1}^B \rho(\|\mathbf{s}_i - p_j\|; \phi, \kappa) \\ \frac{1}{\|A\|^2} \iint \rho(\|\mathbf{s} - \mathbf{s}'\|; \phi, \kappa) ds ds' &\approx \frac{1}{B^2} \sum_{j=1}^B \sum_{j'=1}^B \rho(\|p_j - p_{j'}\|; \phi, \kappa) \end{aligned}$$

These approximations can be made arbitrarily accurate by increasing the number of sample points B .

It is important to note that although our estimate for τ_u^2 is for a specific county, we use data from all available monitors (including ones outside the county) to obtain the estimate of τ_u^2 .

3.4 Risk estimation

Let $\mathbf{w}_t = (w(\mathbf{s}_1, t), \dots, w(\mathbf{s}_n, t))$ represent the observed data for all of the monitors on day t . Using the spatial-temporal model and the joint distribution in (5), the conditional distribution of the true unobserved county-wide average x_t given the data is

$$x_t \mid \mathbf{w}_t \sim \mathcal{N}(\mu_{x,t} + H'_{12}H_{11}^{-1}(\mathbf{w}_t - \boldsymbol{\mu}_t), \sigma^2(H_{22} - H'_{12}H_{11}^{-1}H_{12})) \quad (6)$$

where maximum likelihood estimates of σ^2 , ϕ , and κ are plugged in where necessary. We use this conditional distribution to adjust estimates from health risk models in two different ways.

3.4.1 Two-stage Bayesian model

The main approach we describe for adjusting risk estimates for spatial misalignment error is a two-stage Bayesian model. In the first stage we estimate $r(x_t \mid \mathbf{w}_t)$, which is simply the distribution in (6), i.e. the posterior distribution of x_t given the data \mathbf{w}_t for each t . The second stage uses $r(x_t \mid \mathbf{w}_t)$ as an informative prior for x_t and estimates the joint posterior distribution of θ and $\mathbf{x} = (x_1, \dots, x_T)$, given the health data \mathbf{y} and the observed pollutant data \mathbf{w} ,

$$\begin{aligned} p(\theta, \mathbf{x} \mid \mathbf{y}, \mathbf{w}) &\propto p(\mathbf{y} \mid \theta, \mathbf{x}, \mathbf{w})r(\mathbf{x} \mid \mathbf{w})\pi(\theta) \\ &= \left[\prod_{t=1}^T p(y_t \mid \theta, x_t, \mathbf{w}_t) \right] \left[\prod_{t=1}^T r(x_t \mid \mathbf{w}_t) \right] \pi(\theta) \end{aligned}$$

where $\mathbf{y} = (y_1, \dots, y_T)$, $\mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_T)$, and $\pi(\theta)$ is a diffuse prior distribution. The likelihood terms $p(y_t \mid \theta, x_t, \mathbf{w}_t)$ represent the Poisson likelihood used for the time series model relating pollutant exposure to health outcomes. Details on that part of the model are shown in Section 4.2.

Note that in a full Bayesian model the target marginal posterior of θ given the data is

$$\begin{aligned} p(\theta \mid \mathbf{y}, \mathbf{w}) &= \int p(\theta, \mathbf{x} \mid \mathbf{y}, \mathbf{w}) d\mathbf{x} \\ &= \int p(\theta \mid \mathbf{x}, \mathbf{y}, \mathbf{w})p(\mathbf{x} \mid \mathbf{y}, \mathbf{w}) d\mathbf{x} \end{aligned}$$

Our two-stage approach effectively assumes that $p(\mathbf{x} \mid \mathbf{y}, \mathbf{w}) \approx r(\mathbf{x} \mid \mathbf{w})$, thus cutting the feedback between \mathbf{x} and \mathbf{y} . Given that previous studies have indicated that the relationship between the health outcome \mathbf{y} and the pollutant exposure \mathbf{x} is generally weak, this assumption is not likely to be unreasonable. Furthermore, we obtain the tremendous practical advantage of separating the model

into two stages so that the substantial work of fitting the spatial model in the first stage (as well as the model checking) can be conducted separately from the parameter estimation in the second stage.

The posterior distribution of θ given the data can be sampled using Markov Chain Monte Carlo techniques. Specifically, we use a hybrid Gibbs sampler and alternate sampling from the full conditional distributions of \mathbf{x} and θ . Details of the sampling algorithm can be found in the Appendix.

3.4.2 Regression calibration

As an alternative to the two-stage Bayesian model, we can use a regression calibration type of approach (Carroll et al., 2006). Using the distribution in (6), we can calculate $\mathbb{E}[x_t | \mathbf{w}_t]$ for each day t . Then, substituting $\mathbb{E}[x_t | \mathbf{w}_t]$ in place of \bar{w}_t and conducting the standard analysis described in Section 2 would give us the regression calibrated estimate of our risk parameter θ . This method should produce estimates that are similar to the two-stage model and has the advantage that it requires substantially less computation

4 Application

Daily counts of hospital admissions for the period 2000–2006 were obtained from billing claims of enrollees in the US Medicare system. Each billing claim contains the date of service, disease classification (International Classification of Diseases 9th Revision [ICD-9] codes), age, and county of residence. We considered as an outcome urgent or emergency hospital admissions for cardiovascular diseases, which were calculated using ICD-9 codes (Dominici et al., 2006). The daily counts of hospitalizations were calculated by summing the hospital admissions for each disease of interest recorded as a primary diagnosis. To calculate daily hospitalization rates, we constructed a parallel time series of the numbers of individuals enrolled in Medicare that were at risk in each county on each day. We restricted the analysis to the 20 largest counties in the country with at least 100 observations on components of $\text{PM}_{2.5}$ over the 7 year period of 2000–2006.

In the US, chemical components of $\text{PM}_{2.5}$ are typically measured once every six days and patterns of missing data vary depending on when monitors began collecting data regularly. For this analysis

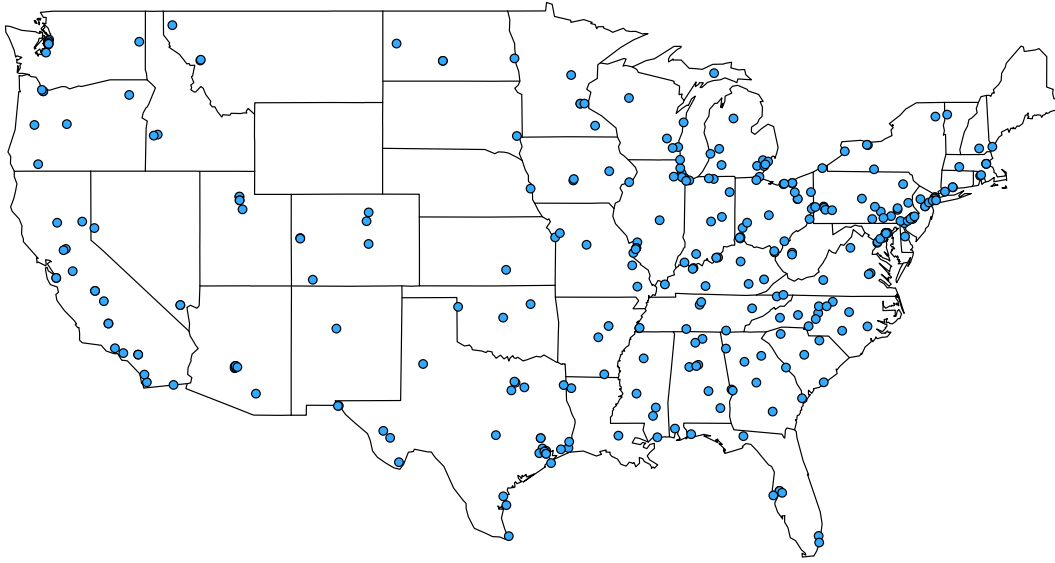


Figure 1: Locations of 313 chemical speciation monitors in the US, 2000–2006.

we do not attempt to impute the data and only use the days for which we have observations for all relevant variables. Our analysis was limited to the components making up a large fraction of the total $PM_{2.5}$ mass or co-varying with total mass. These components were sulfate, nitrate, silicon, elemental carbon (EC), organic carbon matter (OCM), sodium ion, and ammonium ion. These seven components, in aggregate, constituted 83% of the total $PM_{2.5}$ mass, whereas all other components individually contributed less than one percent (Bell et al., 2006). In total, we obtained data from 313 chemical speciation monitors across the US. A map of the monitor locations is shown in Figure 1. National temperature and dew point temperature data were obtained from the National Climatic Data Center on the Earth-Info CD database.

4.1 Estimation of spatial–temporal model

For each of the seven chemical components, we fit the spatial-temporal model described in Section 3. Because we are only interested in short-term associations between chemical components and health outcomes, before fitting the model, we detrended the time series for each component, removing any seasonal fluctuations and long-term trends. To detrend the data, we fit a linear model with the monitor value for the component as the response and the day of the week, a natural spline of time

with 49 degrees of freedom (i.e. 7 degrees of freedom per year), and temperature as predictors. The residuals from this model were then used as our new chemical component predictor variable.

In Table 1 we show the estimates of the parameters in the Matérn model for each of the seven components. Asymptotic standard errors for the parameters were obtained by inverting the Hessian matrix estimated from maximizing the log-likelihood. The parameter estimates produce correlation

	σ	ϕ	κ
Sulfate	1.25 _(0.0016)	5.32 _(0.0716)	0.12 _(0.0012)
Nitrate	1.23 _(0.0016)	2.86 _(0.0500)	0.16 _(0.0022)
Silicon	0.39 _(0.0005)	6.57 _(0.1296)	0.06 _(0.0009)
EC	0.61 _(0.0008)	8.98 _(0.2794)	0.03 _(0.0006)
OCM	1.54 _(0.0020)	7.14 _(0.1574)	0.06 _(0.0009)
Sodium Ion	0.49 _(0.0006)	10.25 _(0.6736)	0.01 _(0.0005)
Ammonium	0.88 _(0.0011)	4.78 _(0.0709)	0.12 _(0.0013)

Table 1: Maximum likelihood parameter estimates for Matérn model with asymptotic standard errors in parentheses.

functions that generally agree with our knowledge of the spatial distribution of these chemical components. For sulfate, nitrate, and ammonium the decrease in correlation with distance is generally slower than for silicon, EC, OCM, and sodium ion. For all of the components, the small estimates of κ indicate a non-differentiable process and produce a rapid decrease in correlation at short distances followed by a slower decrease at longer distances.

4.1.1 Model checking

To examine the fit of the spatial-temporal model we divided the n monitors randomly into eight groups and conducted an 8-fold cross validation. At each iteration, we held-out one group of monitors and fit the model using the monitors from the remaining seven groups. We then used the fitted model to predict all the values at the held-out monitors. We used mean squared error to summarize the model's performance.

	Sulfate	Nitrate	Silicon	EC	OCM	Sodium Ion	Ammonium
RMSE	1.00	0.47	0.04	0.09	0.85	0.02	0.40
RMSE / Median	0.34	0.50	0.68	0.17	0.26	0.34	0.31

Table 2: Root mean squared error for prediction of spatial-temporal model at held-out monitors.

Table 2 shows root mean squared errors for the spatial-temporal model from the 8-fold cross validation. We also show the RMSE divided by the median levels of each chemical component so that the RMSE can be compared across the different scales of variation of the chemical components. Table 2 is meant to give some sense of the prediction accuracy of the spatial-temporal model. However, it should be noted that the ultimate purpose of the model is to predict the county-wide average chemical component level. Predicting chemical component concentrations at specific locations is used here as a measure of model fit, albeit an imperfect one.

We also checked the assumptions of the measurement error model in (4), which assumes that the errors u_t are independent of the true county-wide average values x_t . For each county in the analysis we held out the monitors inside the county and fit the spatial-temporal model to the remaining monitors. We then compared the estimate of the county-wide average based on the monitor values inside the county, \bar{w}_t , and the posterior mean of the true county-wide average, x_t , obtained from the model. After calculating $u_t = \bar{w}_t - x_t$, the correlation between x_t and u_t , on average across locations and chemical components, was 0.27, indicating a relatively weak correspondence between the two.

4.1.2 Spatial misalignment error

Using the fitted spatial-temporal model, we can compute estimates of τ_u^2 , the spatial misalignment error variance from (4), for each of the seven chemical components and each county. For each component we can also calculate estimates of σ_x^2 , the temporal variance of the true unobserved county-wide average component level. Table 3 shows for each of the 20 counties and seven components the spatial misalignment error variance ratio, which is the ratio of the spatial misalignment error variance to the variance of true county-wide average, i.e. τ_u^2/σ_x^2 . We can see that the components silicon, EC, OCM, and sodium ion generally have much larger error ratios than the other three components. In

addition, there appears to be a correspondence between the error ratios and size of the county (by area) as well as the number of monitors.



	Area (km ²)	# Monitors	Sulfate	Nitrate	Silicon	EC	OCM	Sodium Ion	Ammonium
Los Angeles, CA	10,518	1	1.08	1.08	2.39	4.95	2.40	13.92	1.21
Cook, IL	2,449	4	0.22	0.20	0.50	1.05	0.49	3.00	0.23
Maricopa, AZ	23,836	5	0.68	0.84	1.08	1.66	1.05	3.76	0.74
San Diego, CA	10,878	2	0.59	0.60	1.24	2.57	1.25	7.13	0.64
Queens, NY	283	1	0.53	0.42	1.36	3.08	1.38	9.15	0.56
Dallas, TX	2,278	2	0.41	0.38	0.93	2.00	0.95	5.70	0.45
Wayne, MI	1,591	3	0.28	0.29	0.64	1.34	0.64	3.76	0.33
King, WA	5,506	6	0.62	0.71	0.88	1.29	0.86	2.74	0.64
Santa Clara, CA	3,343	2	0.65	0.70	1.30	2.49	1.34	6.48	0.71
Broward, FL	3,122	1	0.89	0.83	1.99	4.15	2.01	11.93	0.96
Riverside, CA	18,667	1	1.93	2.12	3.53	6.63	3.52	17.34	2.08
New York, NY	59	1	0.49	0.38	1.18	2.68	1.19	7.84	0.49
Philadelphia, PA	350	3	0.17	0.14	0.45	1.07	0.47	3.12	0.18
Cuyahoga, OH	1,187	2	0.47	0.45	1.00	2.06	1.00	5.57	0.50
Clark, NV	20,488	2	1.04	1.12	1.94	3.52	1.94	8.61	1.15
Bronx, NY	109	3	0.20	0.17	0.44	0.95	0.44	2.67	0.20
Allegheny, PA	1,891	3	0.29	0.26	0.64	1.34	0.63	3.75	0.32
Sacramento, CA	2,501	2	0.46	0.42	0.99	2.12	1.04	5.89	0.51
Hennepin, MN	1,442	2	0.67	0.66	1.30	2.35	1.26	6.00	0.71
Franklin, OH	1,398	1	0.71	0.64	1.71	3.73	1.73	10.85	0.78

Table 3: Area, number of monitors, and ratios τ_u^2/σ_x^2 for the seven chemical components and 20 US counties, 2000–2006.

4.1.3 Monitor coverage

To assess the relationship between the number of monitors in a county and the degree of spatial misalignment error, we fit a linear regression of the form

$$\log_2 \frac{\tau_u^2}{\sigma_x^2} = \beta_0 + \beta_1 \log_2(\# \text{ of Monitors in a County}) + \beta_2 \log_2(\text{County Area})$$

The data for this model are taken from Table 3 and we fit a separate model for each chemical component. In this model, β_1 can be interpreted as the change in the spatial misalignment error ratio with a doubling of the number of monitors in the county. In order for this quantity to be interpretable we have to adjust for the area of the county first.

In each of the panels of Figure 2 we plot for each county the partial residuals for the number of monitors on the x -axis and the partial residuals for the spatial misalignment error ratio τ_u^2/σ_x^2 on the y -axis (both are on a \log_2 scale). The estimated values of β_1 for each component are shown inside the individual panels in Figure 2. Given $\hat{\beta}_1$, we can compute $100 \times (2^{\hat{\beta}_1} - 1)$, which is the percent change in the proportion of spatial misalignment error associated with doubling the number of monitors in a county, adjusting for a county's area.

It appears that for each component, counties with more monitors in them (adjusted for the total area of the county) have smaller spatial misalignment error variance ratios. For example, when estimating county-wide average nitrate levels, counties with 2 monitors rather than 1 have an approximately 35% decrease in the spatial misalignment error variance ratio. For elemental carbon, the benefit of going from 1 to 2 monitors is a 46% decrease in the error ratio and for sodium ion there is a 48% decrease. While the spatial misalignment error generally decreases for all PM components when the number of monitors increases, the benefit is especially pronounced for silicon, EC, OCM, and sodium ion. Thus, for more spatially heterogeneous components, more benefit (i.e. less spatial misalignment error) is gained by additional monitor coverage than for components that are spatially homogeneous.

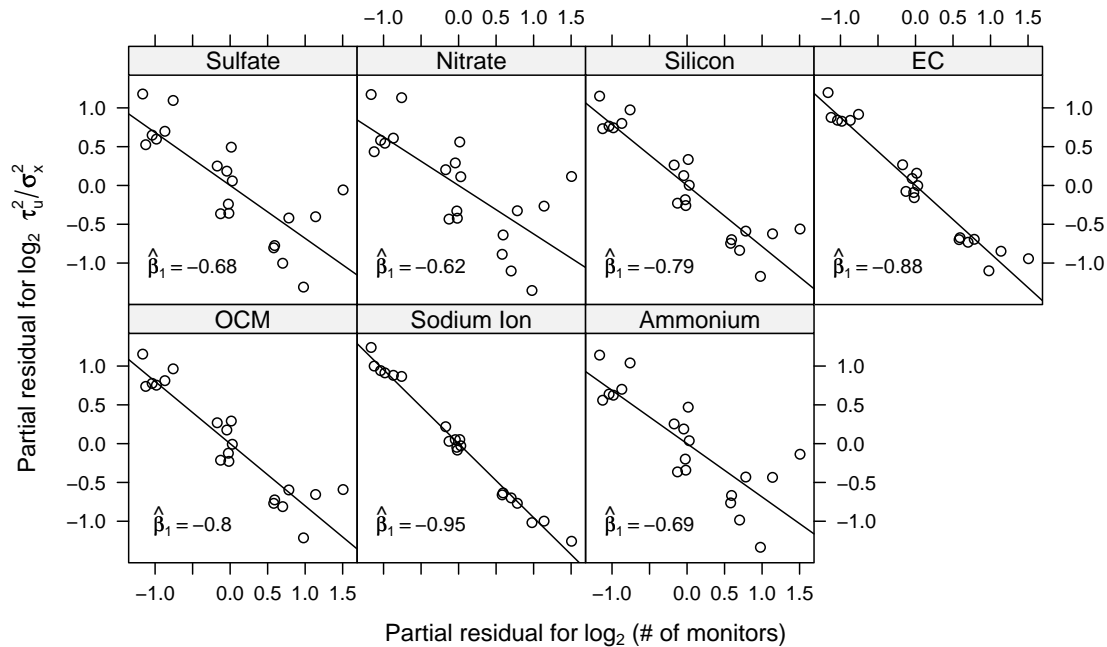


Figure 2: The (\log_2) spatial misalignment error ratio versus the (\log_2) number of monitors in a county, adjusted for county area.

4.2 Risk estimation

For each county, we fit the following log-linear Poisson model to the health and chemical component data, which is an extended version of the model in (1).

$$\begin{aligned} \log \mathbb{E}[Y_t] = & \alpha + \theta x_t + s(t, \lambda_1) + s(\text{temp}_t, \lambda_2) + s(\text{dewpt}_t, \lambda_3) \\ & + \log(\# \text{ at risk}_t) + \gamma \text{dow}_t + s(\text{temp}_{1-3,t}, \lambda_4) + s(\text{dewpt}_{1-3,t}, \lambda_5) \end{aligned}$$

Y_t is the number of admissions for cardiovascular disease and x_t is the county-wide average of the chemical component being examined. When we use the regression calibration approach, x_t is estimated using the regression calibration function $\mathbb{E}[x_t | \mathbf{w}_t]$ and with the two-stage Bayesian model the values of x_t are sampled from the full conditional distribution within the MCMC iterations. We assume that the variables other than the pollutant variable are measured without error.

The additional terms in the model represent an offset for the number of people at risk, the day of the week, a smooth function of the three day running mean of temperature, and a smooth function

of the three day running mean of dew point temperature. The running means of temperature are included to capture the effects of temperature in the winter (Samet et al., 1998). For the smooth functions we used $\lambda_1 = 49$, $\lambda_2 = \lambda_4 = 6$, and $\lambda_3 = \lambda_5 = 3$. These values have been used previously and generally capture the variation in season as well as temperature (Peng et al., 2006; Welty and Zeger, 2005).

Estimates of the risk parameter θ are shown in Table 4. The results are shown for estimates obtained using standard maximum likelihood, which ignores the spatial misalignment problem, as well as regression calibration and the two-stage Bayesian model which explicitly adjust for spatial misalignment error. For the Bayesian model we use the posterior mean as our point estimate and the posterior standard deviation as our measure of uncertainty. Table 4 shows the five largest counties (by population) of the 20 counties we examined.

The point estimates from the regression calibration procedure and the two-stage Bayesian model are generally in agreement given the uncertainties. The standard errors for the Bayesian estimates tend to be slightly larger than the standard errors obtained using the regression calibration procedure. We see in Table 4 that the chemical components that exhibited greater spatial misalignment error resulted in larger adjustment for their estimated risk parameters. However, for components where there did not appear to be a strong association to begin with (e.g. silicon), the adjusted estimates from the regression calibration and Bayesian models were not substantially different from the maximum likelihood estimates.

Figures 3 and 4 show the risk estimates from the three models for elemental carbon and sulfate, respectively, in all 20 counties. Elemental carbon is relatively heterogeneous spatially and we see that for a number of counties the adjustment for spatial misalignment changes the point estimate substantially. Sulfate is much more spatially homogeneous and we can see from Figure 4 that the adjustment for spatial misalignment has little effect here. Analogous figures for the remaining components can be found in the supplementary materials.

Method	Sulfate	Nitrate	Silicon	EC	OCM	Sodium Ion	Ammonium	
Los Angeles, CA	MLE	0.66 _(1.10)	0.10 _(0.24)	-0.38 _(0.59)	1.25 _(0.61)	-0.65 _(0.88)	0.34 _(0.41)	0.17 _(0.44)
	RegCal	1.29 _(2.17)	-0.02 _(0.42)	-0.47 _(1.15)	3.74 _(2.02)	-2.90 _(1.73)	4.01 _(2.35)	-0.02 _(0.80)
	Bayes	1.30 _(2.22)	0.00 _(0.41)	0.81 _(1.56)	4.02 _(1.99)	-2.64 _(1.61)	4.93 _(1.64)	-0.04 _(0.82)
Cook, IL	MLE	-0.79 _(0.65)	-0.16 _(0.42)	0.31 _(0.30)	-0.06 _(0.56)	-0.79 _(0.87)	2.05 _(0.93)	-0.55 _(0.55)
	RegCal	-0.78 _(0.77)	0.32 _(0.52)	0.50 _(0.46)	0.25 _(1.24)	-0.05 _(1.29)	5.85 _(3.09)	-0.26 _(0.68)
	Bayes	-0.80 _(0.76)	0.30 _(0.52)	0.63 _(0.49)	0.02 _(1.33)	-0.27 _(1.32)	4.13 _(0.98)	-0.27 _(0.68)
Maricopa, AZ	MLE	-0.49 _(6.93)	0.65 _(1.42)	0.18 _(0.28)	0.66 _(0.71)	0.02 _(1.10)	-1.74 _(1.00)	3.79 _(3.46)
	RegCal	9.80 _(10.67)	2.76 _(2.72)	0.46 _(0.82)	3.21 _(2.73)	1.28 _(2.68)	-1.40 _(5.56)	6.91 _(4.83)
	Bayes	8.96 _(11.03)	1.78 _(2.89)	0.40 _(0.86)	2.89 _(2.85)	1.00 _(2.75)	1.36 _(4.30)	6.49 _(5.16)
San Diego, CA	MLE	-0.16 _(2.39)	0.67 _(0.67)	1.15 _(0.90)	2.00 _(1.24)	3.02 _(1.27)	0.05 _(0.41)	0.10 _(1.15)
	RegCal	0.58 _(3.58)	0.50 _(0.76)	1.05 _(1.15)	4.13 _(2.51)	2.46 _(2.11)	-1.30 _(2.80)	-0.06 _(1.33)
	Bayes	-0.38 _(3.76)	0.64 _(0.78)	1.46 _(1.12)	5.97 _(2.70)	3.03 _(2.21)	-2.02 _(2.73)	-0.26 _(1.38)
Queens, NY	MLE	2.31 _(1.03)	0.64 _(0.96)	0.96 _(0.87)	0.53 _(0.78)	0.17 _(0.96)	0.10 _(0.06)	1.44 _(0.88)
	RegCal	2.47 _(1.11)	0.90 _(1.00)	0.84 _(1.37)	2.15 _(1.28)	0.80 _(1.19)	1.24 _(0.68)	1.82 _(0.98)
	Bayes	2.68 _(1.15)	0.82 _(1.04)	1.04 _(1.55)	1.35 _(1.43)	0.63 _(1.22)	1.27 _(0.68)	1.90 _(1.04)

Table 4: Estimates and standard errors (in parentheses) of the percent increase in cardiovascular hospital admissions associated with a 1 interquartile range increase in the chemical component using standard maximum likelihood (“MLE”), regression calibration (“RegCal”), and the two-stage Bayesian model (“Bayes”). The values shown are $100 \times (\exp(\hat{\theta} \times \text{IQR}) - 1)$. For the two-stage Bayesian model, $\hat{\theta}$ is the posterior mean. The IQRs for each component are expressed in $\mu\text{g}/\text{m}^3$ and are 3.06 (Sulfate), 1.64 (Nitrate), 0.07 (Silicon), 0.40 (EC), 3.18 (OCM), 0.11 (Sodium Ion), and 1.35 (Ammonium).

Elemental Carbon

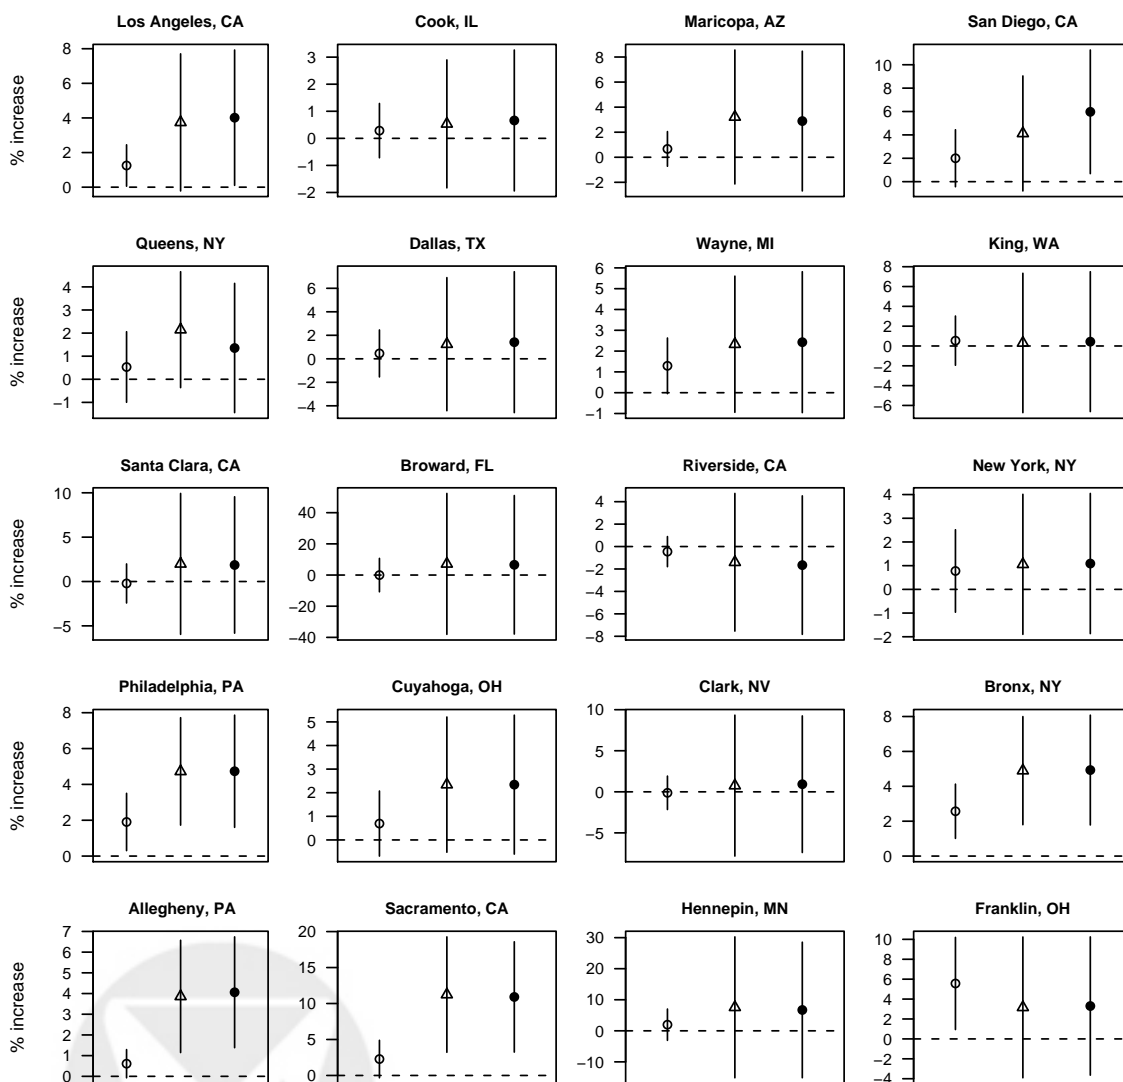


Figure 3: Percent increase in cardiovascular hospital admissions per 1 interquartile range ($0.40 \mu\text{g}/\text{m}^3$) increase in elemental carbon estimated using maximum likelihood (open circle), regression calibration (triangle), and the two-stage Bayesian model (filled circle).




Figure 4: Percent increase in cardiovascular hospital admissions per 1 interquartile range ($3.06 \mu\text{g}/\text{m}^3$) increase in elemental carbon estimated using maximum likelihood (open circle), regression calibration (triangle), and the two-stage Bayesian model (filled circle).

5 Discussion

We have presented a statistical model for estimating and adjusting for spatial misalignment error in time series studies of air pollution and health. We demonstrated a regression calibration approach that is computationally very efficient as well as a two-stage Bayesian model. These approaches were used to estimate the risk of cardiovascular hospitalization associated with exposure to chemical components of particulate matter.

Our findings indicate that the degree of spatial misalignment is a function of monitor coverage within a county and the pollutant of interest. In general, a decrease in the amount of area covered per monitor in a county is associated with lower spatial misalignment error, however this effect is far more pronounced for pollutants that are heterogeneous such as sodium ion, silicon, and EC. In particular, for counties with only one monitor, it would seem there may be some benefit to using one additional monitor in the county.

However we must be cautious in translating our results into specific recommendations for monitor placement. Our analysis did not take into account the placement of the monitors in the county and the density of the population. For counties with very high population density concentrated in one specific area, a single monitor might be sufficient and it would make little sense to place monitors throughout the county where few people live. Further research is certainly needed in order to determine the optimal usage and placement of pollution monitors. The design of monitoring networks is generally important for regulatory purposes and cost tradeoffs may be prohibitive as well. Our methodology could be extended to provide a quantitative basis for assessing the placement of monitors.

In our application we found that health risk estimates for EC, a pollutant exhibiting large spatial heterogeneity, were generally larger using approaches that accounted for spatial misalignment. The adjusted estimates exhibited the classic bias-variance tradeoff with their substantially inflated standard errors. This increased statistical uncertainty comes from accounting for the lack of information about levels of the pollutant at all points in the county. For homogeneous pollutants (e.g. sulfate, ammonium), the information available from the monitoring network provides sufficient information about county-wide levels so that adjusted risk estimates are largely unchanged.

Both the regression calibration approach and the two-stage Bayesian model produced similar adjusted risk estimates in the 20 counties analyzed. However, in general, regression calibration should be used with care in generalized linear models with nonlinear link functions because a separate bias can be introduced in those situations (Carroll et al., 2006). Although we employed log-linear models in this application, the small size of the regression coefficients likely produced a nearly linear model, potentially explaining the similar estimates given by the two approaches.

There exist several alternative approaches to estimate exposures for locations without monitors and also for temporal periods without measurement for use in health-based air pollution research. These include inverse distance weighting incorporating population density (Ivy et al., 2008), air quality modeling (Bell, 2006), and kriging (Leem et al., 2006). However, some of these approaches are computational intensive and rely on datasets beyond those typically used in epidemiological settings. Our approach has the advantage of relying exclusively on existing datasets that are commonly used in air pollution studies, and our results can provide insight into the interpretation of results based on simple county-average exposure metrics, such as results across different components or areas with different monitor coverage.

One improvement that could be made to our model is to incorporate population density information if available. In our spatial-temporal pollutant process, we create the county-wide averages by integrating the process against a uniform density over the county boundary. However, if we could specify a function $g(\mathbf{s})$ which indicates the population density at location \mathbf{s} , then we could compute the following county-wide average instead,

$$x_t = \int w(\mathbf{s}, t)g(\mathbf{s}) ds.$$

This value x_t might reflect more accurately the population-level exposure than our current approach. We consider this an interesting avenue for future work. Another natural extension of our model would be to extend it to the multi-site setting where data are available for many locations. Our approach here was to apply the single-location model described in Section 3 independently to each available county. However, a unified model for multiple locations would estimate location-specific risks simultaneously while taking into account the spatial correlation between the county-wide average pollutant levels.

There are many sources of measurement error in the analysis of air pollution and health data and much previous work has focused on the mismatch between personal and ambient exposures to an airborne pollutant (Dominici et al., 2000; Zeger et al., 2000). This is indeed an important problem, although it is typically not one that can be dealt with usefully given the types of data that are routinely available. A key assumption made in previous time series analyses of air pollution and health data is that the pollutant of interest is spatially homogeneous and that the spatial misalignment error is negligible (Zeger et al., 2000). Often, this assumption was justifiable, given that most previous analyses focused on pollutants such as the total mass of particulate matter (PM₁₀, PM_{2.5}), ozone, other pollutants which have been shown to be fairly spatially homogeneous over relatively long distances (Samet et al., 2000b).

Ultimately, the best way to address the problem of spatial misalignment might be to move away from the county-based summaries of the outcome of interest when possible and begin using summaries with finer spatial resolution, such as zip codes. Unfortunately, many types of health data are simply not available at finer spatial resolution and we often must accept what is available. Further, due to activity patterns, a high spatial resolution does not necessarily better capture personal exposure than a larger area when individuals move between areas (e.g., live in one zip code, but work in another). Thus, there is a strong need for methods that address spatial misalignment of air pollutant concentrations used in health studies. In such cases the methods proposed here should be useful for determining the magnitude of the errors incurred and for obtaining adjusted risk estimates.

Funding

This work was supported in part by the United States Environmental Protection Agency through STAR grant to the Johns Hopkins University [RD832417]. It has not been subjected to the Agency's required peer and policy review and therefore does not necessarily reflect the views of the Agency and no official endorsement should be inferred. Funding was also provided by the National Institute for Environmental Health Sciences [ES012054-03 to R.P.]; and by the National Institute for Environmental Health Sciences Center in Urban Environmental Health [P30ES03819 to R.P.].

Collection of Biostatistics
Research Archive

Acknowledgments

The authors thank Keita Ebisu for assistance with the chemical components database and Francesca Dominici for helpful comments on the manuscript.

References

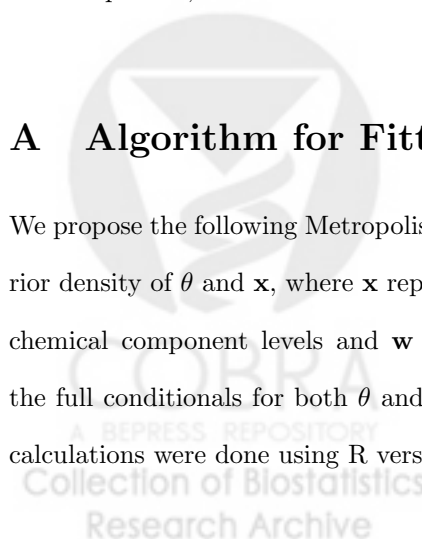
- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2004), *Hierarchical Modeling and Analysis for Spatial Data*, Chapman & Hall/CRC.
- Bell, M. L. (2006), “The use of ambient air quality modeling to estimate individual and population exposure for human health research: a case study of ozone in the Northern Georgia region of the United States,” *Environment International*, 32, 586–593.
- Bell, M. L., Dominici, F., Ebisu, K., Zeger, S. L., and Samet, J. M. (2006), “Spatial and temporal variation in PM_{2.5} chemical composition in the United States for health effects studies,” Tech. Rep. 118, Johns Hopkins University, <http://www.bepress.com/jhubiostat/paper118>.
- Carroll, R. J., Ruppert, D., Stefanski, L., and Crainiceanu, C. M. (2006), *Measurement Error in Nonlinear Models: A Modern Perspective*, Chapman & Hall/CRC.
- Cheng, G. and Kosorok, M. R. (2008), “The penalized profile sampler,” *Journal of Multivariate Analysis*, to appear, doi:10.1016/j.jmva.2008.05.001.
- Dominici, F., Peng, R. D., Bell, M. L., Pham, L., McDermott, A., Zeger, S. L., and Samet, J. M. (2006), “Fine Particulate Air Pollution and Hospital Admission for Cardiovascular and Respiratory Diseases,” *Journal of the American Medical Association*, 295, 1127–1134.
- Dominici, F., Zeger, S. L., and Samet, J. M. (2000), “A measurement error correction model for time-series studies of air pollution and mortality,” *Biostatistics*, 2, 157–175.
- Fuentes, M., Song, H.-R., Ghosh, S. K., Holland, D. M., and Davis, J. M. (2006), “Spatial Association between Speciated Fine Particles and Mortality,” *Biometrics*, 62, 855–863.

- Gryparis, A., Paciorek, C. J., Zeka, A., Schwartz, J., and Coull, B. A. (2008), “Measurement error caused by spatial misalignment in environmental epidemiology,” *Biostatistics*, to appear.
- Ivy, D., Mulholland, J. A., and Russell, A. G. (2008), “Development of ambient air quality population-weighted metrics for use in time-series health studies,” *Journal of Air and Waste Management Association*, 58, 711–720.
- Jones, G. L., Haran, M., Caffo, B. S., and Neath, R. (2006), “Fixed-Width Output Analysis for Markov Chain Monte Carlo,” *Journal of the American Statistical Association*, 101, 1537–1547.
- Katsouyanni, K., Toulomi, G., Samoli, E., Gryparis, A., LeTertre, A., Monopoli, Y., Rossi, G., Zmirou, D., Ballester, F., Boumghar, A., and Anderson, H. R. (2001), “Confounding and Effect Modification in the Short-term Effects of Ambient Particles on Total Mortality: Results from 29 European Cities within the APHEA2 Project,” *Epidemiology*, 12, 521–531.
- Leem, J. H., Kaplan, B. M., Shim, Y. K., Pohl, H. R., Gotway, C. A., Bullard, S. M., Rogers, J. F., Smith, M. M., and Tylanda, C. A. (2006), “Exposures to air pollutants during pregnancy and preterm delivery,” *Environmental Health Perspectives*, 114, 905–910.
- Nocedal, J. and Wright, S. J. (1999), *Numerical Optimization*, Springer.
- Peng, R. D., Chang, H. H., Bell, M. L., McDermott, A., Zeger, S. L., Samet, J. M., and Dominici, F. (2008), “Coarse particulate matter air pollution and hospital admissions for cardiovascular and respiratory diseases among Medicare patients,” *Journal of the American Medical Association*, 299, 2172–2179.
- Peng, R. D., Dominici, F., and Louis, T. A. (2006), “Model choice in time series studies of air pollution and mortality (with discussion),” *Journal of the Royal Statistical Society, Series A*, 169, 179–203.
- R Development Core Team (2008), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.
- Ribeiro, P. J. and Diggle, P. J. (2001), “geoR: a package for geostatistical analysis,” *R News*, 1, 14–18, ISSN 1609-3631.

- Samet, J., Zeger, S., Kelsall, J., Xu, J., and Kalkstein, L. (1998), “Does Weather Confound or Modify the Association of Particulate Air Pollution with Mortality?” *Environmental Research, Section A*, 77, 9–19.
- Samet, J. M., Dominici, F., Zeger, S. L., Schwartz, J., and Dockery, D. W. (2000a), *The National Morbidity, Mortality, and Air Pollution Study, Part I: Methods and Methodological Issues*, Health Effects Institute, Cambridge MA.
- Samet, J. M., Zeger, S. L., Dominici, F., and et al. (2000b), *The National Morbidity, Mortality, and Air Pollution Study, Part II: Morbidity and Mortality from Air Pollution in the United States*, Health Effects Institute, Cambridge MA.
- Touloumi, G., Atkinson, R., Le Tertre, A., Samoli, E., Schwartz, J., Schindler, C., Vonk, J., Rossi, G., Saez, M., Rabszenko, D., and Katsouyanni, K. (2004), “Analysis of health outcome time series data in epidemiological studies,” *Environmetrics*, 15, 101–117.
- Welty, L. J. and Zeger, S. L. (2005), “Are the Acute Effects of PM₁₀ on Mortality in NMMAPS the Result of Inadequate Control for Weather and Season? A Sensitivity Analysis using Flexible Distributed Lag Models.” *American Journal of Epidemiology*, 162, 80–88.
- Zeger, S. L., Thomas, D., Dominici, F., Samet, J. M., Schwartz, J., Dockery, D., and Cohen, A. (2000), “Exposure measurement error in time-series studies of air pollution: Concepts and consequences,” *Environmental Health Perspectives*, 108, 419–426.

A Algorithm for Fitting the Two-Stage Bayesian Model

We propose the following Metropolis-Hastings sampling algorithm for sampling from the joint posterior density of θ and \mathbf{x} , where \mathbf{x} represents the vector of unobserved true county-wide daily average chemical component levels and \mathbf{w} is the vector observed daily monitor average values. Briefly, the full conditionals for both θ and \mathbf{x} are sampled using a Metropolis-Hastings rejection step. All calculations were done using R version 2.7.1 (R Development Core Team, 2008).



We make use of the profile likelihood for θ and \mathbf{x} which profiles out the many nuisance parameters in the likelihood (Cheng and Kosorok, 2008). These parameters include the spline coefficients for the smooth function of time and the nonlinear functions of temperature and dew point temperature. Let $\boldsymbol{\eta}$ be the vector of nuisance parameters in the full likelihood. We evaluate the profile likelihood $L_p(\theta, \mathbf{x}) = \max_{\boldsymbol{\eta}} L_f(\theta, \mathbf{x}, \boldsymbol{\eta})$ where for each pair of values (θ, \mathbf{x}) we maximize the full Poisson likelihood L_f with respect to $\boldsymbol{\eta}$. This can be done simply by fitting a standard generalized linear model with an offset for θ and \mathbf{x} . Then in the Metropolis-Hastings steps to sample θ and \mathbf{x} , we use the profile likelihood to calculate the acceptance ratios.

1. *Sampling θ .* We use a random walk Metropolis-Hastings step so that the proposal distribution at step i is

$$\theta^* \mid \theta^{(i-1)} \sim \mathcal{N}(\theta^{(i-1)}, \widehat{V})$$

where \widehat{V} is the variance of the maximum likelihood estimate of θ . The acceptance ratio is then calculated as

$$r = \frac{L_p(\theta^*, \mathbf{x}^{(i-1)})\pi(\theta^*)}{L_p(\theta^{(i-1)}, \mathbf{x}^{(i-1)})\pi(\theta^{(i-1)})}$$

where the prior $\pi(\theta)$ is taken to be a Normal distribution with mean 0 and standard deviation 10.

2. *Sampling \mathbf{x} .* We sample the vector \mathbf{x} as a block of length T , where T is the number of observations we have for a given county. The prior distribution for \mathbf{x} is the posterior distribution specified in equation (6) and derived from fitting the spatial model described in Section 3. For the proposal distribution, we use the distribution in (6) so that the proposal equals the prior. Given a proposal value \mathbf{x}^* , we compute the acceptance ratio, which in this case is simply the profile likelihood ratio,

$$r = \frac{L_p(\theta^{(i-1)}, \mathbf{x}^*)}{L_p(\theta^{(i-1)}, \mathbf{x}^{(i-1)})}$$

Each sampler was run for 10,000 iterations. Convergence of the chains was diagnosed by estimating Monte Carlo standard errors of the parameters using the method of batch means described in Jones et al. (2006).

SUPPLEMENTARY MATERIALS



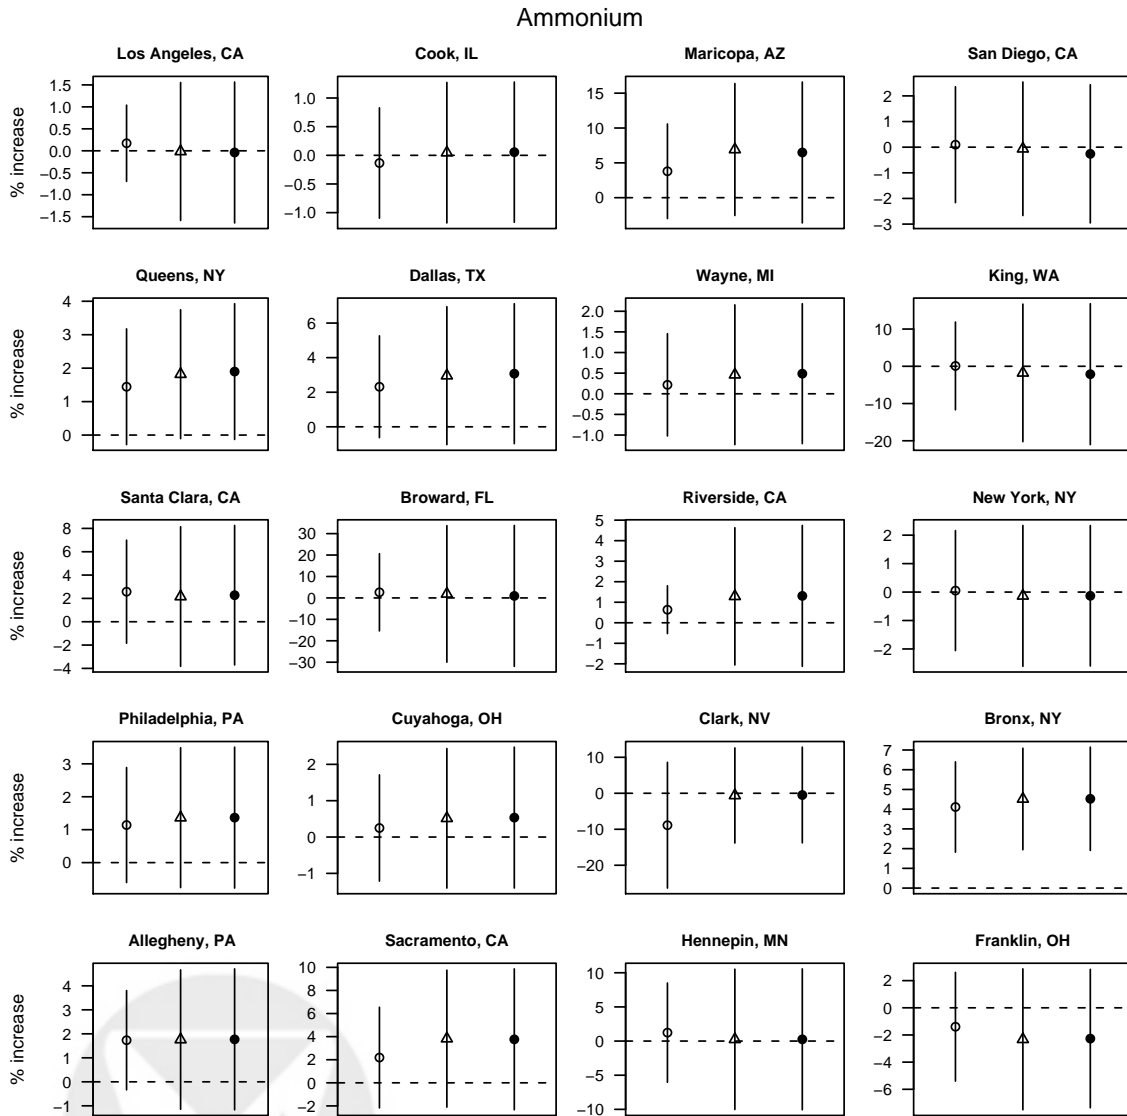


Figure 5: Percent increase in cardiovascular hospital admissions per 1 interquartile range ($1.35 \mu\text{g}/\text{m}^3$) increase in ammonium estimated using maximum likelihood (open circle), regression calibration (triangle), and the two-stage Bayesian model (filled circle).




Figure 6: Percent increase in cardiovascular hospital admissions per 1 interquartile range ($1.64 \mu\text{g}/\text{m}^3$) increase in nitrate estimated using maximum likelihood (open circle), regression calibration (triangle), and the two-stage Bayesian model (filled circle).

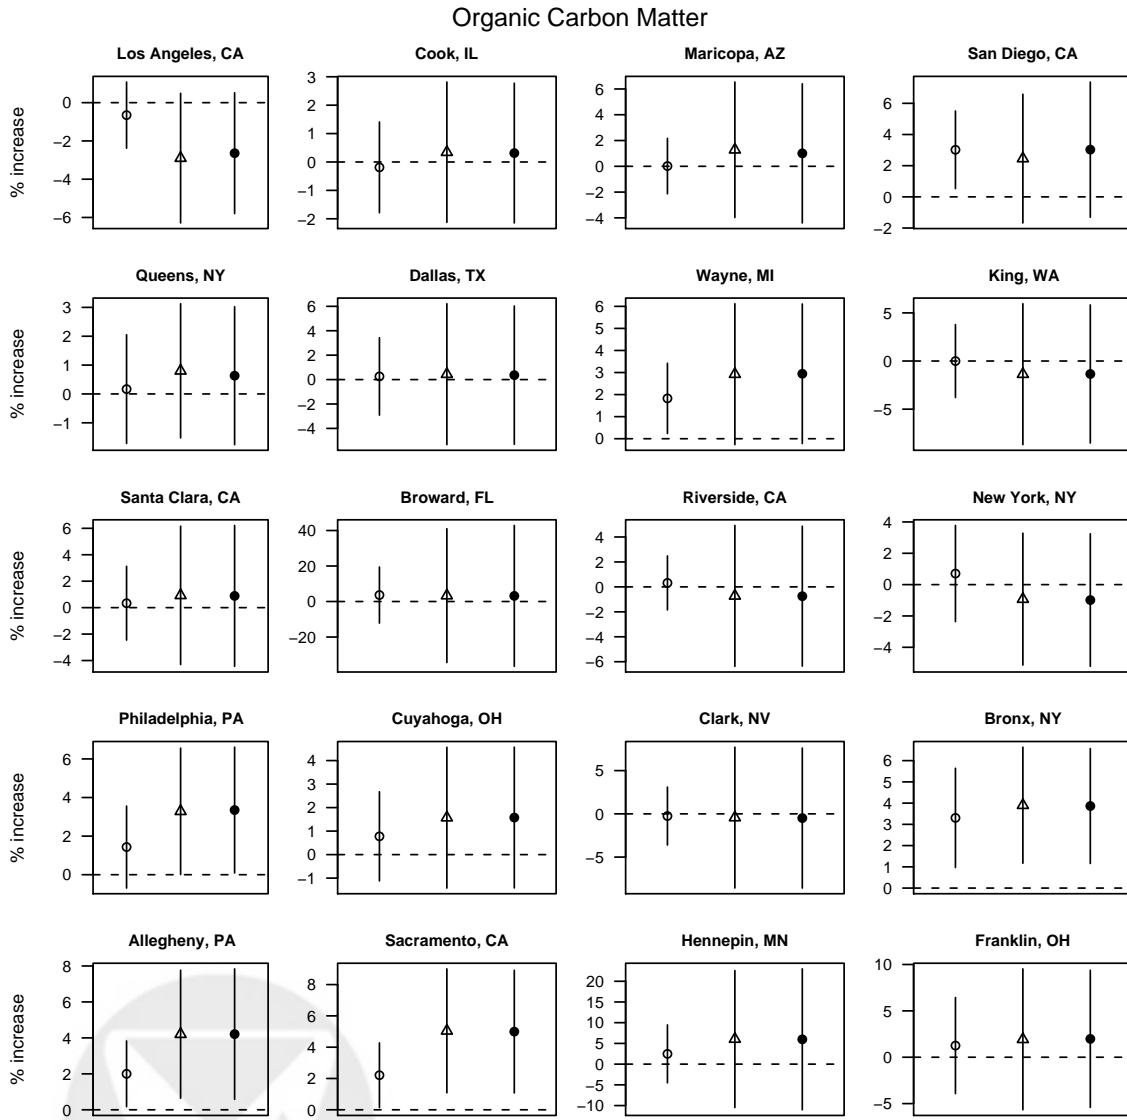


Figure 7: Percent increase in cardiovascular hospital admissions per 1 interquartile range ($3.18 \mu\text{g}/\text{m}^3$) increase in organic carbon matter estimated using maximum likelihood (open circle), regression calibration (triangle), and the two-stage Bayesian model (filled circle).




Figure 8: Percent increase in cardiovascular hospital admissions per 1 interquartile range (0.11 μg/m³) increase in sodium ion estimated using maximum likelihood (open circle), regression calibration (triangle), and the two-stage Bayesian model (filled circle).




Figure 9: Percent increase in cardiovascular hospital admissions per 1 interquartile range ($0.07 \mu\text{g}/\text{m}^3$) increase in silicon estimated using maximum likelihood (open circle), regression calibration (triangle), and the two-stage Bayesian model (filled circle).

