6-24-2009

A Spatio-Temporal Approach for Estimating Chronic Effects of Air Pollution

Sonja Greven
Johns Hopkins University, sgreven@jhsph.edu

Francesca Dominici
Johns Hopkins University

Scott L. Zeger
Johns Hopkins University

A Spatio-Temporal Approach for Estimating Chronic Effects of Air Pollution

Sonja Greven Francesca Dominici Scott Zeger
Department of Biostatistics, Johns Hopkins University, Baltimore, MD 21205.
E-mail: sgreven@jhsph.edu

Abstract

Estimating the health risks associated with air pollution exposure is of great importance in public health. In air pollution epidemiology, two study designs have been used mainly. Time series studies estimate acute risk associated with short-term exposure. They compare day-to-day variation of pollution concentrations and mortality rates, and have been criticized for potential confounding by time-varying covariates. Cohort studies estimate chronic effects associated with long-term exposure. They compare long-term average pollution concentrations and time-to-death across cities, and have been criticized for potential confounding by individual risk factors or city-level characteristics.

We propose a new study design and a statistical model, which use spatio-temporal information to estimate the long-term effects of air pollution exposure on life expectancy. Our approach avoids confounding by time-varying covariates and individual or city-level risk factors. By estimating the increase in life expectancy due to decreases in long-term air pollution concentrations, it provides easily interpretable values for public policy purposes. We develop a suitable backfitting algorithm that permits efficient fitting of our model to large spatio-temporal data sets. We evaluate spatio-temporal correlation in the data and obtain appropriate standard errors. We apply our methods to the Medicare Cohort Air Pollution

Funding was provided by the U.S. Environmental Protection Agency (EPA) (grants RD-83241701 and RD-83362201). Although the research described in this article has been funded wholly or in part by the U.S. EPA, it has not been subjected to the agency's required peer and policy review and does not necessarily reflect the views of the agency, and no official endorsement should be inferred.

Study, including data on fine particulate matter ($PM_{2.5}$) and mortality for 18.2 million Medicare enrollees from 814 locations in the U.S. during an average of 65 months in 2000-2006. Supplemental material including R code implementing our methods is provided in a web appendix.

Keywords: Backfitting Algorithm, Environmental Epidemiology, Particulate Matter, Spatio-temporal Modeling

1. INTRODUCTION

The Clean Air Act (Environmental Protection Agency, last amended in 1990) requires the U.S. Environmental Protection Agency (EPA) to set National Ambient Air Quality Standards for seven pollutants considered harmful. Air quality standards for several air pollutants have since also been adopted by the European Union. Implementation of these standards led to decreases in air pollution concentrations in the United States (Bachmann, 2008). From a public policy and public health perspective, it is of importance to assess whether these decreases have also led to an improvement in morbidity and mortality for the general population (Health Effects Institute, 2003). Standards are reviewed periodically, with evidence from epidemiologic studies playing a large role in the public policy process (Kaiser, 1997; Greenbaum et al., 2001; Samet et al., 2003).

Two types of study design have been mainly used to estimate the association between air pollution and mortality.

Time series studies (see for example Schwartz and Dockery, 1992; Spix et al., 1993; Kelsall et al., 1997) estimate acute effects of short-term exposure to air pollutants. They compare day-to-day variations in mortality with those in air pollution concentrations. Multi-site time series studies (Katsouyanni et al., 1997; Samet et al., 2000; Samoli et al., 2008; Wong et al., 2008) combine the evidence and statistical uncertainty across geographical locations, providing an overall estimate of the association that accounts for the within-location and across-location variance (Dominici et al., 2000; Dominici, 2002). Due to the focus on short-term effects, time series studies do not allow an assessment of the years of life-time lost due to air pollution (Künzli et al., 2001). Potential confounders in time series studies are time-varying variables such as weather or seasonal effects, as well as slowly varying unmeasured factors. Typically, smooth functions of weather variables and calendar time are included in the regression model to account for temporal confound-

ing. However, results have been found to be sensitive to the flexibility granted to these smooth functions (Samoli et al., 2001; Klemm and Mason, 2003; Dominici et al., 2004; Peng et al., 2006), and time series studies have sometimes been criticized with regard to potential residual confounding (Vedal, 1997; Lumley and Sheppard, 2000; Moolgavkar, 2005).

Cohort studies (see for example Dockery et al., 1993; Pope et al., 2002; Laden et al., 2006; Eftim et al., 2008) estimate the chronic effects of long-term exposure to air pollution. They compare across locations long-term average air pollution concentrations and time-to-death in cohorts. Cohort studies allow the estimation of life expectancy lost due to air pollution (Künzli et al., 2001; Rabl, 2003). They have been criticized (Moolgavkar, 1994; Vedal, 1997; Gamble, 1998; Moolgavkar, 2005), due to the difficulty of fully accounting for all potential confounders, including individual risk factors and location-specific characteristics.

In this paper, we propose a new study design, statistical model, and estimation procedure for estimating the long-term effects of air pollution. Our approach can be denoted as a spatio-temporal cohort study, because it makes use of all the available spatio-temporal variation in the data. It thus differs from time-series studies, which make use of temporal variation in the data, as well as cohort studies, which use spatial variation. As in cohort studies, we evaluate chronic effects of long-term exposure to air pollution, and the resulting years of life-time lost. Here, long-term exposure is defined as the average concentration over the previous year, although longer (moving) time windows can be chosen.

Starting from a proportional hazard model, we derive a Poisson regression model and estimate two regression coefficients. The first coefficient (the global association) estimates the association between the national trend in the air pollutant and the national trend in mortality rates, and is likely to be confounded by other factors that vary slowly on the national level, such as smoking prevalence. The second coefficient (the local association), estimates the association between deviations of location-specific trends in the air pollutant from the national trend, and deviations of location-specific trends in mortality rates from the national trend. We show that the local association is not affected by confounding due to unmeasured time-varying variables, individual-level risk factors or location-level characteristics.

We derive a backfitting algorithm that makes use of the specific model structure to obtain an efficient implementation of our approach. This enables the fitting of our model

to very large spatio-temporal data sets. We evaluate spatio-temporal correlation in the data and derive appropriate standard errors for the estimates of regression coefficients and associated increases in life expectancy. We apply our methods to the Medicare Cohort Air Pollution Study (MCAPS), which includes individual-level information on time of death, age, gender and race on a population of 18.2 million Medicare enrollees from 814 locations in the U.S. for the period 2000-2006. We report new evidence on fine particulate matter ($PM_{2.5}$) and life expectancy.

We first introduce the data and our statistical model in Sections 2.1-2.3, proposing a decomposition of the available spatio-temporal information to avoid confounding. The backfitting algorithm for fitting our proposed regression model is described in Section 2.4. In Sections 2.5-2.6, we explore the spatio-temporal correlation in the data and derive appropriate standard errors for estimates of regression coefficients and associated increases in life expectancy. In Section 3., we apply our methods to a population of 18.2 Medicare enrollees from the MCAPS. Section 4. concludes with a discussion. A web appendix provides all R code used to implement the methods and produce the results in this paper.

2. METHODS

2.1 The Medicare Cohort Air Pollution Study Data

We construct a retrospective cohort study, by linking ambient levels of $PM_{2.5}$ to mortality data by monitor during the period 2000-2006 (see also Zeger et al., 2008, for details).

Specifically, we obtain data from 1,006 $PM_{2.5}$ monitors for the period 2000-2006 from the EPA monitoring network (`http://www.epa.gov/oar/data/`). In our analysis, we include data from 814 monitors in the continental U.S. which have measurements for at least four calendar years with no less than ten months of four daily values or more each. We divide the country into three geographical regions. These are the eastern region, the central region from the Mississippi River to the Sierra Nevada range, and the western United States (Zeger et al., 2008). Monitor locations and regional affiliation are depicted in Figure 1.

[Figure 1 about here.]

We define long-term exposure as the average of daily $PM_{2.5}$ levels over the previous year. These yearly averages are calculated as follows. First, to fill small gaps in the

data, we smooth the $PM_{2.5}$ time series at each location using a linear regression with the daily $PM_{2.5}$ values as the response, and with thin plate splines of time with four degrees of freedom per year as the predictor. For gaps longer than 90 days, we smooth the $PM_{2.5}$ time series before and after the gap separately. Second, for each month, we calculate yearly averages of $PM_{2.5}$ using the 365 predicted daily values from this model up to and including the respective month. In case of missing values, 350 days are deemed sufficient to compute the yearly average. The 814 monitors provide up to 70 monthly measurements of yearly average $PM_{2.5}$ concentrations from December 2000 to September 2006. Summary statistics are given in Table 1.

We then link $PM_{2.5}$ data to the mortality data as follows: the same $PM_{2.5}$ exposure from a given monitoring site is assigned to all Medicare enrollees residing in a ZIP code with a geographic centroid within a six mile radius from that site. The Medicare data provides demographic information (age, gender, race), and individual-level information on survival, with time of death or censoring precise up to the month. The data set includes about 18.2 million enrollees and 3.2 million deaths in total, with an average of 10.4 million people enrolled in the cohort in any given month (Table 1).

[Table 1 about here.]

2.2 The Statistical Model

First, we specify the following proportional hazards model

$$h^c(a,t) \;=\; h^c(a)\exp(PM_t^c\beta), \tag{1}$$

where $h^c(a,t)$ denotes the hazard of dying at age a and time t for location c, and $h^c(a)$ is a location-specific baseline hazard function. PM_t^c is the average of the $PM_{2.5}$ levels at location c over the 12 months prior and including time t.

While the variables age a and time t are continuous variables in principle, the information in the Medicare data on time point of death or censoring is only precise up to the month. We thus discretize the time domain as follows. We measure t in monthly intervals, and denote the set of months with observations for location c by $\mathcal{T}_c$, where $c = 1, \ldots, C$. In a given month t, we count deaths and time at risk as belonging to monthly age interval a, $a = 1, \ldots, A$, if the contributing person turned 65 in month $t-a$. Assuming a constant hazard within each monthly age interval leads to a piecewise exponential survival model for life-tables (see Holford, 1976) for each location.

With a study population of 18.2 million, fitting model (1) is computationally unfeasible. Instead, we use the log-linear regression model

$$\log \mathrm{E}(Y_{at}^c) \quad = \quad \log(T_{at}^c) + \log(h^c(a)) + PM_t^c\beta, \tag{2}$$

with the assumption that each Y_{at}^c is an independent (across calendar time, space and age-months) Poisson variable, conditional on T_{at}^c and PM_t^c. Here, Y_{at}^c is the number of deaths at age-month a in month t for location c, and T_{at}^c is the total time subjects of age a at location c were at risk of dying during month t. Absent the exact time of death or dropout, we approximate T_{at}^c by N_{at}^c, defined as the number of Medicare enrollees of age a with a ZIP code of residence in location c at the beginning of the month. In the appendix, using results by Holford (1980) and Laird and Olivier (1981), we show that under the piecewise exponential survival model, model (1) is equivalent with regard to likelihood-based inference to model (2). Independence assumptions that are made for this equivalence are independence between different locations and birth-month cohorts for model (1), and independence between locations c, months t and age-months a for model (2). We evaluate the justification of these independence assumptions in Section 2.5.

To make computation feasible and avoid excessive zero cell counts, we further assume a constant hazard of dying over one-year age intervals and after age 90. This allows us to collapse ages a into one-year intervals, and to combine all ages over 90 into one age group. Each of the resulting 1.4 million observations $(Y_{at}^c, N_{at}^c, PM_t^c)$ then describes the mortality rate among people being a years of age at location c during month t, with average $\mathrm{PM}_{2.5}$ exposure PM_t^c during the previous year. For each location c, we model the log-hazard function $\log(h^c(a))$ in (2) using thin-plate splines of age with three degrees of freedom, plus a location-specific intercept. An additional indicator for ages over 90 allows for a potential discontinuity in the hazard function due to the mixture of hazard values in this last group.

In model (2), β denotes the increase in the log-hazard of dying in a given month for an increase of 1 $\mu\mathrm{g/m}^3$ in average $\mathrm{PM}_{2.5}$ concentrations during the previous year. Note that differences between locations in average $\mathrm{PM}_{2.5}$ levels across the entire study period are accounted for by the location-specific log-hazard function $\log(h^c(a))$. Only changes over time in $\mathrm{PM}_{2.5}$ concentrations contribute to the estimation of β, avoiding cross-sectional confounding by individual-level risk factors or location-level characteristics.

2.3 Using spatio-temporal information to control for confounding

Absent confounding and measurement error, model (2) allows estimation of the effect of long-term exposure to $PM_{2.5}$ on life expectancy. However, confounding is a common problem in air pollution studies (see Table 4).

To address confounding, we propose to rewrite model (2) as follows

$$\log \mathrm{E}(Y_{at}^c) \;\; = \;\; \log(N_{at}^c) + \log(h^c(a)) + (PM_t^c - \overline{PM_t})\beta_1 + \overline{PM_t}\beta_2, \qquad (3)$$

where $\overline{PM_t}$ denotes the average of the yearly $PM_{2.5}$ averages in month t across locations. We estimate $\overline{PM_t}$ from the 540 monitors with complete time-series, to avoid spurious trends induced by differing monitors contributing over time.

In model (3), β_1 and β_2 both measure the strength of the association between $PM_{2.5}$ and mortality, but use different types of information.

More specifically, the parameter β_2 provides evidence as to whether nationally, $PM_{2.5}$ and mortality rates are decreasing over time in parallel across the study period. The parameter β_2 can be interpreted as the increase in the national log-mortality rate in a given month and age group, for an increase by 1 μg/m^3 in the national average $PM_{2.5}$ concentration during the previous year. Estimation of this national or "global" parameter β_2 can be confounded by variables which vary slowly nationally, in a similar fashion as both $PM_{2.5}$ and mortality. Potential confounders are national trends in the economy and/or national trends in smoking prevalence. These national trends can be associated with the national trends in air pollution and mortality, and thus can bias the estimation of β_2.

By contrast, the parameter β_1 measures the strength of the evidence that mortality rates decline faster (slower) than the national average in locations where $PM_{2.5}$ levels also decline faster (slower) than the national average. In other words, β_1 measures the association between deviations of the local $PM_{2.5}$ trend from the national $PM_{2.5}$ trend and deviations of the local trend in mortality rates from the national trend in mortality rates, adjusted by the association between the national trends in $PM_{2.5}$ and mortality rates.

Estimation of this "local" parameter β_1 can be affected by a different type of confounding than estimation of β_2. Estimation of β_1 can be biased by confounders that change faster than their national average in some locations, and slower than their national average in others, if the pattern of where these fast and slow changes occur is similar to those of both $PM_{2.5}$ and mortality. Consider the example of the variable "health consciousness". Health consciousness is not a confounder for estimation of β_1 if it differs by location

across the country, or if it increases nationally. However, health consciousness could be a confounder if 1) some local communities developed a heightened health consciousness during the six years of the study period, while others did not (or to a lesser extent), and 2) the locations with changes stronger than the national average and weaker than the national average tended to be the ones where these changes were also stronger/weaker than the national average for $PM_{2.5}$ and mortality. For example, increases in health consciousness might be related to decreases in $PM_{2.5}$ concentrations in a community, but also to decreases in the prevalences of smoking and unhealthy diet, and thus to decreases in mortality rates. These considerations illustrate that while it is possible for estimation of β_1 to be confounded, confounding is less likely to occur than for β_2, or more generally for regression coefficients from time series or cohort studies. In summary, by design confounding by unmeasured time-varying covariates and by individual-level or area-level risk factors cannot bias the estimation of β_1.

An illustration of the interpretation of β_1 and β_2, and of potential confounders is given in the web appendix. Absent confounding, β_1 and β_2 are equal and can be collapsed into the single "overall" coefficient β in (2). Separate estimation of β_1 and β_2 in (3) therefore can help to diagnose unmeasured confounding (Janes et al., 2007).

2.4 Estimation Using the Backfitting Algorithm

Fitting model (3) directly is computationally very demanding. First, this is due to the high dimensionality of the data set with 1.4 million observations $(Y_{at}^c, N_{at}^c, PM_t^c)$, where a ranges through 26 (mostly yearly) age groups, t through the on average 65 months per location, and c through the 814 locations. Second, this is due to the complexity of the model, which specifies a log-hazard function $\log(h^c(a))$ with 5 degrees of freedom for each of the 814 locations.

To reduce the dimensionality of the problem, we use a backfitting algorithm (Buja et al., 1989). We first center the two variables $(PM_t^c - \overline{PM_t})$ and $\overline{PM_t}$ by location by subtracting their location-mean values. This improves orthogonality with $\log(h^c(a))$, but does not change the interpretation or value of β_1 and β_2, as any effect of location-mean $PM_{2.5}$ levels is absorbed by the location-specific intercept. The backfitting algorithm then proceeds as follows.

- Initialize $\beta_1^{(0)} = \beta_2^{(0)} = 0$ and $\log(h^c(a))^{(0)} \equiv 0, c = 1, \ldots, C$.

- *Step A:* For iteration j, set the offset to

$$\text{offset1}_{at}^{c\ (j)} = \log(N_{at}^c) + \log(h^c(a))^{(j-1)}$$

for all a, t and c. Fit the Poisson model

$$\log \mathrm{E}(Y_{at}^c) = \text{offset1}_{at}^{c\ (j)} + (PM_t^c - \overline{PM_t})\beta_1 + \overline{PM_t}\beta_2$$

and set $\beta_1^{(j)}$ and $\beta_2^{(j)}$ to the estimated coefficients.

- *Step B:* For iteration j, set the offset to

$$\text{offset2}_{at}^{c\ (j)} = \log(N_{at}^c) + (PM_t^c - \overline{PM_t})\beta_1^{(j)} + \overline{PM_t}\beta_2^{(j)}$$

for all a, t and c. For $c = 1, \ldots, C$, fit the Poisson model

$$\log \mathrm{E}(Y_{at}^c) = \text{offset2}_{at}^{c\ (j)} + \log(h^c(a))$$

to data from location c, and set $\log(h^c(a))^{(j)}$ to the log-hazard function estimated from this model.

- While the change in $\beta_1^{(j)}$ or $\beta_2^{(j)}$ is larger than a certain stop criterion, repeat steps A and B. Conclude with step A.

The algorithm greatly reduces computational complexity by estimating the log-hazard function for each location separately. To investigate potential overdispersion, an overdispersion parameter $\phi = \mathrm{Var}(Y_{at}^c)/\mathrm{E}(Y_{at}^c)$ can be included in the last Step A.

This backfitting algorithm is slightly different from the local scoring algorithm typically employed in estimation for generalized additive models (Hastie and Tibshirani, 1990). There, one backfitting algorithm for additive models (inner loop) is carried out at each Newton-Raphson step (outer loop), and convergence results from the backfitting algorithm for additive models (Buja et al., 1989) carry over directly. Here, we carry out a full iteratively reweighted least squares algorithm (inner loop) for each step of the backfitting algorithm (outer loop). However, convergence of $\beta_1^{(j)}$ and $\beta_2^{(j)}$ to the unique maximum likelihood estimates $\widehat{\beta}_1$ and $\widehat{\beta}_2$ is straightforward, and is shown in the appendix.

We choose the stop criterion of the algorithm as a small relative change in the parameter estimates, $\max\{|(\beta_1^{(j)} - \beta_1^{(j-1)})/\beta_1^{(j-1)}|, |(\beta_2^{(j)} - \beta_2^{(j-1)})/\beta_2^{(j-1)}|\} < 10^{-6}$, which is reached within 4-8 iterations for the MCAPS data.

2.5 Variance Estimates and Spatio-Temporal Correlation

Variance estimators for $\widehat{\boldsymbol{\beta}} = (\widehat{\beta}_1, \widehat{\beta}_2)$ that account for uncertainty in the estimation of the log-hazard functions $\log(h^c(a))$ can be obtained from the likelihood using standard asymptotic theory; details are given in the appendix. These model-based variance estimators are obtained under the assumption of independence across time, age groups, and geographical locations. In this section, we investigate the justification of this independence assumption.

We examine averaged empirical variograms over age, time and space. More specifically, we define the empirical variogram over space, averaged over time and age, as follows. The averaged value for a bin of spatial distances $(\delta_1, \delta_2]$ is given by

$$\frac{1}{N_{\delta_1,\delta_2}} \sum_{c=1}^{C} \sum_{u:\delta_{uc}\in(\delta_1,\delta_2]} \sum_{t\in\mathcal{T}_{cu}} \sum_{a=1}^{A} \frac{1}{2}(r_{at}^c - r_{at}^u)^2,$$

where $r_{at}^c = (y_{at}^c - \widehat{\mu}_{at}^c)/\sqrt{\widehat{\mu_{at}^c}}$ is the standardized residual and $\widehat{\mu}_{at}^c$ the fitted value from model (3) for location c, month t and age a, $\mathcal{T}_{cu}$ is the set of months common to locations c and u, δ_{cu} is the spatial distance between locations c and u, and N_{δ_1,δ_2} is the number of terms in the sum for the interval $(\delta_1, \delta_2]$. As for the usual variogram over space (see for example Diggle and Ribeiro, 2007), the averaged variogram can be compared to the variance estimate $\widehat{\sigma}^2 = \sum_{a,t,c} {r_{at}^c}^2/N$. Complete independence between spatial locations corresponds to variogram values close to $\widehat{\sigma}^2$ for all spatial distances. Averaged variograms over age and time are defined analogously.

2.6 Estimating Years of Life Gained

To evaluate the public health significance of our findings, we estimate the years of life gained due to decreases in $PM_{2.5}$ exposure. For a known hazard function $h(a)$, we can calculate the life expectancy of a 65-year old individual for a given exposure x and effect β as

$$LE(x, \beta) = \sum_a a h(a) \exp(x\beta) \prod_{b<a} \big[1 - h(b)\exp(x\beta)\big],$$

where a runs through the monthly ages starting with 65. We set the hazard function to the average estimated hazard function across all locations, $h(a) \equiv \sum_{c=1}^{C} \widehat{h}^c(a)/C$.

To estimate the increase in life expectancy associated with a decrease in the annual average of $PM_{2.5}$ by 10 $\mu g/m^3$, we compute $\Delta LE(\beta) = LE(x - 10, \beta) - LE(x, \beta)$.

$\Delta LE(\beta)$ is the difference between the life expectancy assuming the personal exposure to be constant and equal to x, and the life expectancy assuming the exposure to be 10 μg/m^3 less than x. We choose x as the average of the $PM_{2.5}$ yearly average concentrations during the first year of the study period. Note that this approach estimates the increase in life expectancy after age 65, a lower bound for the overall increase in life expectancy.

We compute $\Delta LE(\widehat{\beta}_1)$ and $\Delta LE(\widehat{\beta})$, and their approximate standard errors using the Delta method. Details are given in the appendix.

3. RESULTS FROM THE MCAPS STUDY

Yearly average $PM_{2.5}$ concentrations have been decreasing during the study period in most of the study locations (Figure 2), with a pronounced drop in $PM_{2.5}$ levels after September of 2001. $PM_{2.5}$ average levels were comparable between East and West in 2001, and higher than in the central region. The West shows the strongest and most consistent decline over time, which might reflect the stricter California ambient air quality standard of 12 μg/m^3 annual average $PM_{2.5}$ that came into effect July 5, 2003 (California Environmental Protection Agency Air Resources Board, `http://www.arb.ca.gov/research/aaqs/aaqs.htm`). The decline in $PM_{2.5}$ concentrations is less pronounced in the East and center, with higher average levels in 2005 after an initial decrease. The national average is dominated by values from the eastern region, which contributes 518 of the 814 monitors in this study.

[Figure 2 about here.]

Monthly age-standardized mortality rates have decreased over the same time period in all regions (Figure 3). Rates are comparable in the East and center, and lower in the West, with mortality rates peaking in the winter months. Mortality rates are age-standardized to the cohort population in October 2003, the middle of the study period. Maps of yearly average $PM_{2.5}$ concentrations and age-standardized mortality rates by location are also given in the web appendix.

[Figure 3 about here.]

In Table 2 and Figure 4, we report results from model (3) on the association between long-term exposure to $PM_{2.5}$ and mortality. Table 2 gives estimated coefficients for the U.S. and each region, as well as their respective standard errors. We report estimated

local and global coefficients β_1 and β_2 from model (3), and also the overall coefficient β from model (2). The corresponding relative risk estimates are depicted in Figure 4.

Estimated overdispersion parameter values for model (3) range from $\widehat{\phi} = 1.01$ to 1.02 across regions, and results shown are based on a Poisson model without overdispersion. Average empirical variograms (shown in the web appendix) give no indication of correlation between observations over either space, time, or age. We therefore report model-based standard errors assuming independence across locations, months and age-groups.

Based on the global coefficient β_2, we find that a 10 μg/m^3 increase in the national average $PM_{2.5}$ concentration over the previous year is associated with a significant 50% increase in the risk of dying in a given month for our Medicare cohort. This estimate reflects that nationally, mortality rates are declining over the study period in parallel with $PM_{2.5}$. Estimates of the risk increase associated with a 10 μg/m^3 increase in regional yearly average $PM_{2.5}$ levels range from 25% to 47% across the three regions. The smallest value is estimated in the West, where the decrease in mortality rates is the smallest, but the decline in average $PM_{2.5}$ concentrations is the largest (Figures 2 and 3).

Estimates of the local coefficient β_1 are smaller than estimates of β_2 by an order of magnitude. Based on β_1, we find that after adjusting for the association between national trends in mortality and $PM_{2.5}$, a 10 μg/m^3 increase in the local average $PM_{2.5}$ concentration over the previous year is associated with a significant 2.3% increase in the risk of dying in a given month for the local Medicare population. Regional estimates of β_1 are heterogeneous. Estimates in the West, center and East indicate a 1.7%, -0.6% and 6.0% risk increase, respectively. Only the estimate in the eastern region is statistically significant.

Estimates of β lie between those of β_1 and β_2, as they are a weighted average of these two estimates (Janes et al., 2007). However, the large differences between the estimated local and global coefficients β_1 and β_2 indicate that they cannot be combined into a single coefficient β. Indeed, a test for homogeneity $H_0 : \beta_1 = \beta_2$ results in p-values smaller than 0.0001 in all regions.

[Table 2 about here.]

[Figure 4 about here.]

Table 3 gives estimates with 95% confidence intervals (CIs) for the increase in life expectancy associated with a reduction in yearly average $PM_{2.5}$ concentrations. Results

based on the global coefficient indicate that a 10 $\mu g/m^3$ reduction in the yearly national average of $PM_{2.5}$ is associated with an increase in life expectancy of 14.8 months (CI 13.8-15.7 months) in the Medicare population. However, local coefficients result in different estimates. We find a 10 $\mu g/m^3$ reduction in the local yearly average of $PM_{2.5}$ to be associated with an increase in life expectancy of 2.2 months (CI 1.0-3.4 months) in the local population of Medicare enrollees, after adjusting for the association between national trends in $PM_{2.5}$ and mortality. This increase in life expectancy is heterogeneous across regions, with no significant increase found in western and central regions, and an increase of about 5.6 months (CI 3.6-7.6 months) found in the eastern region.

[Table 3 about here.]

4. DISCUSSION

We have developed a new study design, statistical model and estimation procedure for estimating the chronic effects of long-term exposure to air pollution, and the resulting years of life-time lost. The strength of our study design is that it makes use of all of the available spatio-temporal variation in the data, and thereby avoids confounding by time-varying covariates, individual-level risk factors, and area-level characteristics. It overcomes some of the limitations of time series studies and cohort studies, which use only the temporal and only the spatial variation, respectively, to estimate the association between air pollution and mortality (see Table 4).

[Table 4 about here.]

Our model accounts for the changing hazard of dying with age, and therefore allows the estimation of life-time lost due to air pollution. Inclusion of age-varying hazard functions also explains a large fraction of the variability in the data and lowers the overdispersion to one, thus improving the ability to detect small air pollution effects. Our regression model estimates two regression coefficients, the "global" and the "local" coefficients. The global coefficient estimates the association between the national trends in $PM_{2.5}$ and mortality. The local coefficient estimates the association between local trends in $PM_{2.5}$ and mortality, adjusting for the association between the national trends. While the global coefficient is likely to be confounded by time-varying covariates that vary on the national

level similarly to $PM_{2.5}$ and mortality, the local coefficient is unaffected from confounding by time-varying covariates, individual risk factors and area-level characteristics by design.

We have applied our methods to the Medicare Cohort Air Pollution Study. We reported new evidence on fine particulate matter ($PM_{2.5}$) and life expectancy, using data on 18.2 million enrollees from 814 U.S. locations during up to 70 months in 2000-2006. We find that nationally, mortality rates are decreasing, by about 4% for each 1 $\mu g/m^3$ decrease in average $PM_{2.5}$ levels during the previous year (global coefficient). Locally, an additional 1 $\mu g/m^3$ decrease in yearly average $PM_{2.5}$, above and beyond the national decline in $PM_{2.5}$ levels, is associated with an additional 0.23% decrease in the local mortality rates, on top of the national reduction in mortality (local coefficient).

The large difference between global and local parameter estimates indicates that confounding affects estimation of the two coefficients differently, and that the two cannot be combined into a single coefficient. Results for the global coefficient are based on the parallel decrease in national mortality rates and national $PM_{2.5}$ levels. They are likely to be confounded by other variables with a temporal trend on the national level, such as changes in weather patterns, decreases in other pollutants, changes in smoking patterns and dietary habits, or changes in the economy. The estimated local coefficient, which is likely to be much less biased by confounding, is smaller than the estimated global coefficient. However, we still find a significant increase in life expectancy associated with decreases in average $PM_{2.5}$ concentrations based on the local coefficient.

We find that the significant local association is driven by the result from the East. Stronger $PM_{2.5}$ effects in the East than in the West have been found before for long-term (Zeger et al., 2008) as well as short-term effects (Dominici et al., 2006; Bell et al., 2008). This geographical heterogeneity could potentially be due to chemical composition of $PM_{2.5}$, which differs geographically (Bell et al., 2007) and leads to heterogeneity in the toxicity of $PM_{2.5}$ (U.S. EPA, 2004; Peng et al., 2009; Bell et al., 2009), but more research is needed.

Few studies have investigated the association between $PM_{2.5}$ and mortality using temporal changes in long-term average $PM_{2.5}$, typically by comparing locations across two time periods. (For earlier results on total suspended particulates in the 70s and 80s, see Chay et al., 2003; Chay and Greenstone, 2003). In a follow-up of the Six City Study, Laden et al. (2006) found a 27% (CI 5-43%) reduction in mortality risk associated with each 10 $\mu g/m^3$ reduction in $PM_{2.5}$ levels from 1974-89 to 1990-98, adjusting for time

period and $PM_{2.5}$ levels in 1974-1989. Pope et al. (2009) found a 0.61 year (CI 0.22-1.0 year) increase in mean life expectancy associated with each 10 $\mu g/m^3$ $PM_{2.5}$ decrease from 1979-1983 to 1999-2000 in 51 metropolitan areas, adjusting for changes in socioeconomic, demographic and smoking variables. Our estimates of local coefficients are smaller in comparison, although the confidence intervals overlap. Possible explanations include the more recent time period of our study. Increasingly stringent air pollution regulations might not only lower $PM_{2.5}$ concentrations, but also change its composition and toxicity, leading to decreases in air pollution effects in the last ten to fifteen years (Dominici et al., 2007; Breitner et al., 2009). Janes et al. (2007) in a previous analysis of the Medicare cohort did not find any association between local trends in mortality and local trends in yearly average $PM_{2.5}$, after adjusting for the association between national trends. Our study includes more than twice as many calendar months, spatial locations and Medicare enrollees, and not aggregating mortality rates over ten year age groups might further improve our sensitivity to small $PM_{2.5}$ effects.

The estimated 2.3% increase in mortality rates per 10 $\mu g/m^3$ increase in yearly average $PM_{2.5}$ is larger than the typical 0.3% to 1.2% estimate from time series studies (Pope and Dockery, 2006; Dominici et al., 2007). This is to be expected, as time series studies capture acute, but not chronic pollutant effects, and thus underestimate the total effect on mortality (Künzli et al., 2001). Our estimate is smaller than the estimates of chronic $PM_{2.5}$ effects from cohort studies, which typically find a 4 to 21% increase in all-cause mortality rates per 10 $\mu g/m^3$ increase in long-term average $PM_{2.5}$ (Pope et al., 2002; Pope and Dockery, 2006; Zeger et al., 2008; Eftim et al., 2008). Possible explanations include residual confounding by individual and area-level risk factors in cohort studies, and the shorter exposure time in our study, which could potentially lead to some underestimation of effects.

The strengths of our approach is the ability to estimate the chronic effects of air pollution on life expectancy, while accounting for any time-varying covariates, individual risk factors or area-level characteristics that could bias the estimation through confounding by design. Our approach makes use of the full spatio-temporal information available in large spatio-temporal data sets. The good temporal and spatial resolution then allows the assessment of heterogeneity of effects, such as across geographical regions. An efficient backfitting algorithm permits the fitting of our model to such large data sets. Our approach could also be used to estimate the chronic effects of other pollutants.

A limitation of our study is the unavailability of information on individual-level or area-level covariates, although such information could be included in our model. Time-varying individual-level or area-level risk factors could potentially confound estimation of the local coefficient β_1, although inclusion of time-varying area-level characteristics did not greatly change effect estimates in Pope et al. (2009). We also do not investigate co-pollutants of $PM_{2.5}$, which might decrease locally in parallel with $PM_{2.5}$. However, short-term associations of particulate matter with mortality have been found to not be significantly changed by inclusion of gaseous pollutants in the model (Schwartz, 2004; Dominici et al., 2005). Our study is limited to individuals over 65 years old; the estimated increase in life-expectancy associated with a $PM_{2.5}$ reduction thus gives a lower bound for the overall increase.

In our study, we use ambient $PM_{2.5}$ measurements from single monitors to measure $PM_{2.5}$ exposure. From a public policy perspective, decreases in ambient pollutant concentrations and associated decreases in mortality are of interest in assessing the impact of air quality regulations. Moreover, studies have shown that $PM_{2.5}$ is relatively homogeneous within a given county (Dominici et al., 2006; Janes et al., 2007), and ambient $PM_{2.5}$ is a strong proxy of personal $PM_{2.5}$ exposure (Sarnat et al., 2006).

To fully utilize the spatio-temporal variation in the data, we use the $PM_{2.5}$ average concentration over the last year as the relevant long-term exposure measurement. This approach could potentially miss longer-term effects or lag periods. However, the effects of long-term average $PM_{2.5}$ and $PM_{2.5}$ levels in the year of death have been found to be similar (Laden et al., 2006), which suggests reversibility of effects within about a year. This is plausible in light of the reversibility of the much larger increase in cardiovascular risk in smokers within about three years (Dobson et al., 1991; McElduff et al., 1998).

While we did not see any spatio-temporal correlation in the residuals in our analysis, it would be of interest in general to develop robust standard errors for the regression coefficients that do not require an independence assumption across time and space. Other relevant extensions of our model are spatially or age-varying coefficients, to further investigate the observed spatial heterogeneity of effects, as well as potential differences in effects across age groups (see e.g. Zeger et al., 2008). More work is needed to allow the fitting of these more complex models to large data sets such as the MCAPS data.

APPENDIX: DERIVATIONS

Equivalence of Proportional Hazards Model and Poisson Model (Section 2.2)

The equivalence of the two models has been noted by Holford (1980) and Laird and Olivier (1981). First, consider the proportional hazards model (1) with month-wise constant hazard function for one location c and one birth-month cohort, which turns 65 in the same month t_0. For this cohort, the hazard of dying is constant in age-month interval a, and equal to $h^c(a, t_0 + a) = h^c(a) \exp(PM^c_{t_0+a}\beta)$. The likelihood contribution from this cohort then is, analogous to Laird and Olivier (1981),

$$\prod_{a=1}^{A} \left[h^c(a, t_0 + a)^{Y^c_{a,t_0+a}} \exp(-h^c(a, t_0 + a) T^c_{a,t_0+a}) \right]^{\mathcal{I}(t_0+a \in \mathcal{T}_c)},$$

where Y^c_{at} is the number of deaths at age-month a in month t for location c, T^c_{at} is the total time subjects of age a at location c were at risk of dying during month t, $\mathcal{T}_c$ is the location-specific set of observed months, and $\mathcal{I}(\cdot)$ denotes the indicator function.

If we assume that observations which are from different locations or different birth-month cohorts are independent, the full likelihood can be written as

$$L_S(\beta; h^1(1), \ldots, h^C(A)) = \prod_{c=1}^{C} \prod_{a=1}^{A} \prod_{t \in \mathcal{T}_c} (h^c(a) \exp(PM^c_t \beta))^{Y^c_{at}} \exp(-h^c(a) \exp(PM^c_t \beta) T^c_{at}).$$

For the log-linear Poisson model (2), under the assumption of independence between Y^c_{at} and $Y^{\tilde{c}}_{\tilde{a}\tilde{t}}$ if $(a, t, c) \neq (\tilde{a}, \tilde{t}, \tilde{c})$, the likelihood is

$$\begin{aligned} & L_P(\beta; h^1(1), \ldots, h^C(A)) \\ = & \prod_{c=1}^{C} \prod_{a=1}^{A} \prod_{t \in \mathcal{T}_c} \frac{(h^c(a) T^c_{at} \exp(PM^c_t \beta))^{Y^c_{at}} \exp(-T^c_{at} h^c(a) \exp(PM^c_t \beta))}{Y^c_{at}!} \\ \propto & \; L_S(\beta; h^1(1), \ldots, h^C(A)). \end{aligned}$$

As the two likelihoods are proportional, the two models are equivalent with regard to likelihood-based inference.

Convergence of the Backfitting Algorithm (Section 2.4)

The log-likelihood $\ell(\boldsymbol{\beta}, \boldsymbol{\gamma})$ is a function of $\boldsymbol{\beta} = (\beta_1, \beta_2)$ and $\boldsymbol{\gamma}$ of length $5C$, which contains an indicator, three spline basis functions in a and an indicator for ages over 90 for each location (see below, Model-based Variance Estimates). The log-likelihood is based on an exponential family density and is strictly concave, as well as bounded above, with $\ell(\boldsymbol{\beta}, \boldsymbol{\gamma}) \to -\infty$ if one of the coordinates goes to $\pm\infty$. Thus, the maximum likelihood

estimator of ($\boldsymbol{\beta}$, $\boldsymbol{\gamma}$) exists and is unique, and there are no other local maximizers of the log-likelihood.

The backfitting algorithm alternates between maximizing $\ell(\boldsymbol{\beta}^{(j)}, \boldsymbol{\gamma})$ over $\boldsymbol{\gamma}$ for fixed $\boldsymbol{\beta}^{(j)}$, and maximizing $\ell(\boldsymbol{\beta}, \boldsymbol{\gamma}^{(j)})$ over $\boldsymbol{\beta}$ for fixed $\boldsymbol{\gamma}^{(j)}$. This corresponds to the Block Coordinate Descent/Ascent Method, which converges to $\arg\max_{(\boldsymbol{\beta},\boldsymbol{\gamma})} \ell(\boldsymbol{\beta}, \boldsymbol{\gamma})$, as the log-likelihood is stricly concave and bounded above (Abatzoglou and O'Donnell, 1982; Tseng, 2001).

Model-based Variance Estimates (Section 2.5)

The log-likelihood for model (3) can be defined as follows:

$$\ell(\boldsymbol{\beta}, \boldsymbol{\gamma}) \;\propto\; \sum_{c=1}^{C} \sum_{t \in \mathcal{T}_c} \sum_{a=1}^{A} \left\{ Y_{at}^{c} (\boldsymbol{x_t^{c\prime}} \boldsymbol{\beta} + \boldsymbol{z_a^{c\prime}} \boldsymbol{\gamma}) - N_{at}^{c} \exp(\boldsymbol{x_t^{c\prime}} \boldsymbol{\beta} + \boldsymbol{z_a^{c\prime}} \boldsymbol{\gamma}) \right\}.$$

Here, $\boldsymbol{x_{at}^{c}} = \boldsymbol{x_t^{c}}$ of length 2 contains the $PM_{2.5}$ variables for time t and location c, $\boldsymbol{\beta} = (\beta_1, \beta_2)$, and $\boldsymbol{z_a^{c\prime}\gamma}$ models the log-hazard functions $\log(h^c(a))$, where $\boldsymbol{z_{at}^{c}} = \boldsymbol{z_a^{c}}$ of length $5C$ contains an indicator, three spline basis functions in a, and an indicator for ages over 90 for each location $\tilde{c}$, $\tilde{c} = 1, \ldots, C$. This log-likelihood is based on the assumption of independence between all pairs Y_{at}^{c} and $Y_{\tilde{a}\tilde{t}}^{\tilde{c}}$ for which $(a, t, c) \neq (\tilde{a}, \tilde{t}, \tilde{c})$.

The corresponding score equation is

$$S(\boldsymbol{\beta}, \boldsymbol{\gamma}) \;=\; \sum_{c=1}^{C} \sum_{t \in \mathcal{T}_c} \sum_{a=1}^{A} (\boldsymbol{x_t^{c\prime}}, \boldsymbol{z_a^{c\prime}})' \left[Y_{at}^{c} - N_{at}^{c} \exp(\boldsymbol{x_t^{c\prime}} \boldsymbol{\beta} + \boldsymbol{z_a^{c\prime}} \boldsymbol{\gamma}) \right] = \boldsymbol{X}'(\boldsymbol{Y} - \boldsymbol{\mu}) = \boldsymbol{0},$$

where vectors $\boldsymbol{Y}$ and $\boldsymbol{\mu} = \mathrm{E}(\boldsymbol{Y})$ of length $N = A \sum_c \mathcal{T}_c$ contain entries Y_{at}^{c} and $N_{at}^{c} \exp(\boldsymbol{x_t^{c\prime}} \boldsymbol{\beta} + \boldsymbol{z_a^{c\prime}} \boldsymbol{\gamma})$, respectively, and $\boldsymbol{X}$ is the $N \times (2 + 5C)$ matrix with rows $(\boldsymbol{x_t^{c\prime}}, \boldsymbol{z_a^{c\prime}})$, $a = 1, \ldots, A$, $t \in \mathcal{T}_c$, $c = 1, \ldots, C$.

The model-based asymptotic covariance matrix for $(\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\gamma}})$ then is (McCullagh and Nelder, 1989)

$$\left[-\,\mathrm{E}\left(\frac{d}{d(\boldsymbol{\beta}, \boldsymbol{\gamma})} S(\boldsymbol{\beta}, \boldsymbol{\gamma}) \right) \right]^{-1} = \left(\boldsymbol{X}' \operatorname{diag}(\boldsymbol{\mu}) \boldsymbol{X} \right)^{-1},$$

where $\operatorname{diag}(\boldsymbol{\mu})$ denotes the diagonal matrix with the entries in $\boldsymbol{\mu}$ on the diagonal. Asymptotics here are for $\sum_{t \in \mathcal{T}_c} N_{at}^{c} \to \infty$ for each a and c, while A and C are fixed, such that the number of parameters in $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ stays constant.

Approximate standard errors for estimated years of life gained (Section 2.6)

The quantity to be estimated is $\Delta LE(\beta) = LE(x - 10, \beta) - LE(x, \beta) =: g(\beta)$. An approximate standard error for $g(\widehat{\beta})$ using the Delta method is

$$\sigma(g(\widehat{\beta})) \approx |g'(\widehat{\beta})|\, \sigma(\widehat{\beta}),$$

where $\sigma(g(\widehat{\beta}))$ and $\sigma(\widehat{\beta})$ are the standard errors of $g(\widehat{\beta})$ and $\widehat{\beta}$, respectively,

$$\begin{aligned} g'(\beta) &= \frac{\partial}{\partial\beta}LE(x-10,\beta) - \frac{\partial}{\partial\beta}LE(x,\beta), \qquad \text{and} \\ \frac{\partial}{\partial\beta}LE(x,\beta) &= \sum_a ah(a)x\exp(x\beta)\left(\prod_{b<a}[1-h(b)\exp(x\beta)]\right)\left(1-\sum_{c<a}\frac{h(c)\exp(x\beta)}{1-h(c)\exp(x\beta)}\right). \end{aligned}$$

These standard errors are for a given baseline hazard function $h(a)$, and do not account for uncertainty in estimating $h(a)$.

SUPPLEMENTAL MATERIAL

Supplemental material is provided as a web appendix in a single zip file. Therein

Web_appendix.pdf contains 1) illustrations of β_1, β_2 and potential confounders 2) maps of yearly average $PM_{2.5}$ concentrations and age-standardized mortality rates by location 3) average empirical variograms as defined in Section 2.5, which illustrate the lack of correlation in the residuals over space, time or age for the MCAPS data.

R files are provided which were used to implement the methods and produce the results in this paper. More details are given in the text file ReadMe.rtf.

References

Abatzoglou, T. and B. O'Donnell (1982). Minimization by coordinate descent. *Journal of Optimization Theory and Applications 36*(2), 163–174.

Bachmann, J. (2008). Air pollution forecasts and results-oriented tracking. *Air Quality, Atmosphere & Health 1*(4), 203–207.

Bell, M., F. Dominici, K. Ebisu, S. Zeger, and J. Samet (2007). Spatial and temporal variation in $PM_{2.5}$ chemical composition in the United States for health effects studies. *Environmental Health Perspectives 115*(7), 989–995.

Bell, M., K. Ebisu, R. Peng, J. Samet, and F. Dominici (2009). Hospital Admissions and Chemical Composition of Fine Particle Air Pollution. *American Journal of Respiratory and Critical Care Medicine 179*(12), 1115–1120.

Bell, M., K. Ebisu, R. Peng, J. Walker, J. Samet, S. Zeger, and F. Dominici (2008). Seasonal and Regional Short-term Effects of Fine Particles on Hospital Admissions in 202 US Counties, 1999-2005. *American Journal of Epidemiology 168*(11), 1301–1310.

Breitner, S., M. Stölzel, J. Cyrys, M. Pitz, G. Wölke, W. Kreyling, H. Küchenhoff, J. Heinrich, H. Wichmann, and A. Peters (2009). Short-Term Mortality Rates during a Decade of Improved Air Quality in Erfurt, Germany. *Environmental Health Perspectives 117*(3), 448–454.

Buja, A., T. Hastie, and R. Tibshirani (1989). Linear smoothers and additive models. *Annals of Statistics 17*(2), 453–510.

Chay, K., C. Dobkin, and M. Greenstone (2003). The Clean Air Act of 1970 and adult mortality. *Journal of Risk and Uncertainty 27*(3), 279–300.

Chay, K. and M. Greenstone (2003). The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession. *Quarterly Journal of Economics 118*(3), 1121–1167.

Diggle, P. and P. Ribeiro (2007). *Model-based Geostatistics*. Springer.

Dobson, A., H. Alexander, R. Heller, and D. Lloyd (1991). How soon after quitting smoking does risk of heart attack decline? *Journal of Clinical Epidemiology 44*(11), 1247–1253.

Dockery, D. W., A. Pope, X. Xu, J. D. Spengler, J. H. Ware, M. E. Fay, B. G. Ferris, and F. E. Speizer (1993). An association between air pollution and mortality in six U.S. cities. *New England Journal of Medicine 329*(24), 1753–1759.

Dominici, F. (2002). Invited Commentary: Air Pollution and Health - What Can We Learn from a Hierarchical Approach? *American Journal of Epidemiology 155*(1), 11–15.

Dominici, F., A. McDermott, M. Daniels, S. Zeger, and J. Samet (2005). Revised analyses of the National Morbidity, Mortality, and Air Pollution Study: mortality among residents of 90 cities. *Journal of Toxicology and Environmental Health, Part A 68*(13), 1071–1092.

Dominici, F., A. McDermott, and T. Hastie (2004). Improved semiparametric time series models of air pollution and mortality. *Journal of the American Statistical Association 99*(468), 938–948.

Dominici, F., R. Peng, M. Bell, L. Pham, A. McDermott, S. Zeger, and J. Samet (2006). Fine particulate air pollution and hospital admission for cardiovascular and respiratory diseases. *Journal of the American Medical Association 295*(10), 1127–1134.

Dominici, F., R. Peng, S. Zeger, R. White, and J. Samet (2007). Particulate Air Pollution and Mortality in the United States: Did the Risks Change from 1987 to 2000? *American Journal of Epidemiology 166*(8), 880–880.

Dominici, F., J. Samet, and S. Zeger (2000). Combining evidence on air pollution and daily mortality from the 20 largest US cities: a hierarchical modelling strategy. *Journal of the Royal Statistical Society: Series A 163*(3), 263–302.

Eftim, S., J. Samet, H. Janes, A. McDermott, and F. Dominici (2008). Fine Particulate Matter and Mortality: A Comparison of the Six Cities and American Cancer Society Cohorts With a Medicare Cohort. *Epidemiology 19*(2), 209–216.

Gamble, J. (1998). $PM_{2.5}$ and Mortality in Long-term Prospective Cohort Studies: Cause-Effect or Statistical Associations? *Environmental Health Perspectives 106*(9), 535–549.

Greenbaum, D., J. Bachmann, D. Krewski, J. Samet, R. White, and R. Wyzga (2001). Particulate air pollution standards and morbidity and mortality: case study. *American Journal of Epidemiology 154*(12), 78–90.

Hastie, T. and R. Tibshirani (1990). *Generalized Additive Models*. Chapman & Hall, London.

Health Effects Institute (2003). Assessing the Health Impact of Air Quality Regulations: Concepts and Methods for Accountability Research. Communication 11.

Holford, T. (1976). Life tables with concomitant information. *Biometrics 32*(3), 587–597.

Holford, T. (1980). The analysis of rates and of survivorship using log-linear models. *Biometrics 36*(2), 299–305.

Janes, H., F. Dominici, and S. Zeger (2007). Trends in air pollution and mortality - an approach to the assessment of unmeasured confounding. *Epidemiology 18*(4), 416–423.

Kaiser, J. (1997). Showdown over clean air science. *Science 277*(5325), 466–469.

Katsouyanni, K., G. Touloumi, C. Spix, J. Schwartz, F. Balducci, S. Medina, G. Rossi, B. Wojtyniak, J. Sunyer, L. Bacharova, et al. (1997). Short term effects of ambient sulphur dioxide and particulate matter on mortality in 12 European cities: results from time series data from the APHEA project. *British Medical Journal 314*(7095), 1658–1662.

Kelsall, J., J. Samet, S. Zeger, and J. Xu (1997). Air pollution and mortality in Philadelphia, 1974-1988. *American Journal of Epidemiology 146*(9), 750–762.

Klemm, R. and R. Mason (2003). Replication of Reanalysis of Harvard Six-City Mortality Study. In *Revised Analyses of Time-Series of Air Pollution and Health. Special Report*, pp. 165–172. Health Effects Institute: Boston, MA.

Künzli, N., S. Medina, R. Kaiser, P. Quenel, F. Horak, and M. Studnicka (2001). Assessment of deaths attributable to air pollution: should we use risk estimates based on time series or on cohort studies? *American Journal of Epidemiology 153*(11), 1050–1055.

Laden, F., J. Schwartz, F. Speizer, and D. Dockery (2006). Reduction in fine particulate air pollution and mortality: Extended follow-up of the Harvard Six Cities study. *American Journal of Respiratory and Critical Care Medicine 173*(6), 667–672.

Laird, N. and D. Olivier (1981). Covariance analysis of censored survival data using log-linear analysis techniques. *Journal of the American Statistical Association 76*(374), 231–240.

Lumley, T. and L. Sheppard (2000). Assessing seasonal confounding and model selection bias in air pollution epidemiology using positive and negative control analyses. *Environmetrics 11*(6), 705–717.

McCullagh, P. and J. Nelder (1989). *Generalized linear models* (2nd ed.). Chapman & Hall/CRC.

McElduff, P., A. Dobson, R. Beaglehole, and R. Jackson (1998). Rapid reduction in coronary risk for those who quit cigarette smoking. *Australian and New Zealand Journal of Public Health 22*(7), 787–791.

Moolgavkar, S. (1994). Air pollution and mortality. *New England Journal of Medicine 330*(17), 1237–1238.

Moolgavkar, S. (2005). A Review and Critique of the EPAs Rationale for a Fine Particle Standard. *Regulatory Toxicology and Pharmacology 42*(1), 123–144.

Peng, R., M. Bell, A. Geyh, A. McDermott, S. Zeger, J. Samet, and F. Dominici (2009). Emergency Admissions for Cardiovascular and Respiratory Diseases and the Chemical Composition of Fine Particle Air Pollution. *Environmental Health Perspectives 117*(6), 957–963.

Peng, R., F. Dominici, and T. Louis (2006). Model choice in time series studies of air pollution and mortality. *Journal of the Royal Statistical Society Series A 169*(2), 179–203.

Pope, A., M. Ezzati, and D. W. Dockery (2009). Fine-particulate air pollution and life expectancy in the United States. *New England Journal of Medicine 360*(4), 376–386.

Pope, C., R. Burnett, M. Thun, E. Calle, D. Krewski, K. Ito, and G. Thurston (2002). Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution. *Journal of the American Medical Association 287*(9), 1132–1141.

Pope, C. and D. Dockery (2006). Health effects of fine particulate air pollution: lines that connect. *Journal of the Air & Waste Management Association 56*(6), 709–742.

Rabl, A. (2003). Interpretation of air pollution mortality: number of deaths or years of life lost? *Journal of the Air & Waste Management Association 53*(1), 41–50.

Samet, J., F. Dominici, F. Curriero, I. Coursac, and S. Zeger (2000). Fine particulate air pollution and mortality in 20 US cities, 1987-1994. *New England Journal of Medicine 343*(24), 1742–1749.

Samet, J., F. Dominici, A. McDermott, and S. Zeger (2003). New problems for an old design: time series analyses of air pollution and health. *Epidemiology 14*(1), 11–12.

Samoli, E., R. Peng, T. Ramsay, M. Pipikou, G. Touloumi, F. Dominici, R. Burnett, A. Cohen, D. Krewski, J. Samet, et al. (2008). Acute effects of ambient particulate matter on mortality in Europe and North America: results from the APHENA study. *Environmental Health Perspectives 116*(11), 1480–1486.

Samoli, E., J. Schwartz, B. Wojtyniak, G. Touloumi, C. Spix, F. Balducci, S. Medina, G. Rossi, J. Sunyer, L. Bacharova, et al. (2001). Investigating regional differences in short-term effects of air pollution on daily mortality in the APHEA project: a sensitivity analysis for controlling long-term trends and seasonality. *Environmental Health Perspectives 109*(4), 349–353.

Sarnat, S., B. Coull, J. Schwartz, D. Gold, and H. Suh (2006). Factors affecting the association between ambient concentrations and personal exposures to particles and gases. *Environmental Health Perspectives 114*(5), 649–654.

Schwartz, J. (2004). Is the association of airborne particles with daily deaths confounded by gaseous air pollutants? An approach to control by matching. *Environmental Health Perspectives 112*(5), 557–561.

Schwartz, J. and D. Dockery (1992). Particulate air pollution and daily mortality in Steubenville, Ohio. *American Journal of Epidemiology 135*(1), 12–19.

Spix, C., J. Heinrich, D. Dockery, J. Schwartz, G. Völksch, K. Schwinkowski, C. Cöllen, and H. Wichmann (1993). Air pollution and daily mortality in Erfurt, East Germany, 1980-1989. *Environmental Health Perspectives 101*(6), 518–526.

Tseng, P. (2001). Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications 109*(3), 475–494.

U.S. EPA (2004). Air Quality Criteria for Particulate Matter (Final Report, Oct 2004). U.S. Environmental Protection Agency, Washington, DC, EPA 600/P-99/002aF-bF, 2004.

Vedal, S. (1997). Ambient particles and health: lines that divide. *Journal of the Air & Waste Management Association 47*(5), 551–581.

Wong, C.-M., N. Vichit-Vadakan, H. Kan, Z. Qian, and the PAPA Project Teams (2008). Public Health and Air Pollution in Asia (PAPA): A Multicity Study of Short-Term Effects of Air Pollution on Mortality. *Environmental Health Perspectives 116*(9), 1195–1202.

Zeger, S., F. Dominici, A. McDermott, and J. Samet (2008). Mortality in the Medicare Population and Chronic Exposure to Fine Particulate Air Pollution in Urban Centers (2000–2005). *Environmental Health Perspectives 116*(12), 1614–1619.

Table 1: Number of monitors, number of months with $PM_{2.5}$ data, average $PM_{2.5}$ level, number of Medicare enrollees and number of deaths among Medicare enrollees during the period December 2000 to September 2006. Values are medians among locations, with 25th and 75th percentile given in smaller print.

Region	Monitoring Stations	Months with available $PM_{2.5}$ data	Average $PM_{2.5}$ level [$\mu g/m^3$]	Medicare Enrollees	Deaths
West	96	$_{59}70_{70}$	$_{11.8}13.2_{16.3}$	$_{6164}13289_{37556}$	$_{1099}3028_{7270}$
Center	200	$_{57}70_{70}$	$_{10.3}13.1_{14.9}$	$_{7432}14746_{29073}$	$_{819}1808_{3474}$
East	518	$_{68}70_{70}$	$_{10.8}12.9_{14.5}$	$_{7023}14207_{27688}$	$_{1406}2874_{5451}$
U.S.	814	$_{62}70_{70}$	$_{10.8}13.0_{14.7}$	$_{6957}14502_{29058}$	$_{1220}2539_{5120}$

Table 2: Estimated increase in the log-relative risk of dying in a given month per 1$\mu g/m^3$ increase in average $PM_{2.5}$ concentrations during the previous year. The local coefficient β_1 measures the association between local trends in $PM_{2.5}$ and local trends in mortality rates, adjusting for the respective national trends. The global coefficient β_2 measures the association between the $PM_{2.5}$ national trend and the national trend in mortality rates. The overall coefficient β measures the association between local trends in $PM_{2.5}$ and local trends in mortality, not adjusting for national trends.

Region	Monitors	$100 \times \beta_1$ Estimate$_{(S.E.)}$	$100 \times \beta_2$ Estimate$_{(S.E.)}$	$100 \times \beta$ Estimate$_{(S.E.)}$
West	96	$0.165_{(0.124)}$	$2.249_{(0.113)}$	$1.291_{(0.077)}$
Center	200	$-0.062_{(0.192)}$	$3.849_{(0.303)}$	$1.080_{(0.159)}$
East	518	$0.579_{(0.107)}$	$3.245_{(0.105)}$	$1.929_{(0.075)}$
US	814	$0.223_{(0.064)}$	$4.077_{(0.089)}$	$1.562_{(0.051)}$

Table 3: Estimated increase in life expectancy ΔLE in years for a 10 $\mu g/m^3$ reduction in average yearly $PM_{2.5}$ exposure. Assumptions made in the calculation of ΔLE are given in Section 2.6. Estimates are based on the local coefficient β_1 or on the overall coefficient β. Approximate 95% confidence intervals are derived from the standard errors for $\widehat{\beta}_1$ and $\widehat{\beta}$ using the Delta method.

Region	Monitors	$\Delta LE(\widehat{\beta}_1)_{(95\%CI)}$	$\Delta LE(\widehat{\beta})_{(95\%CI)}$
West	96	$_{-0.06}\,0.14_{0.33}$	$_{0.90}1.02_{1.13}$
Center	200	$_{-0.36}$-$0.05_{0.26}$	$_{0.62}0.87_{1.11}$
East	518	$_{0.30}\,0.47_{0.63}$	$_{1.38}1.49_{1.60}$
US	814	$_{0.08}\,0.18_{0.28}$	$_{1.15}1.23_{1.31}$

Table 4: Study designs for assessing the association between air pollution and mortality

Study design	Exposure	Interpretation of risk	Information	Potential confounders	Example references
Time series	short-term	acute	temporal	time-varying covariates	Katsouyanni et al. (1997) Samet et al. (2000) Samoli et al. (2008) Wong et al. (2008)
Cohort study	long-term	chronic	spatial	individual and area-level risk factors	Dockery et al. (1993) Pope et al. (2002) Laden et al. (2006) Eftim et al. (2008)
Spatio-temporal cohort study	long-term	chronic	spatio-temporal	time-varying individual and area-level risk factors	Janes et al. (2007) Greven et al. (2009)[a]

[a] current article

Figure 1: Locations of 814 EPA $PM_{2.5}$ monitoring sites in the continental United States used for the analysis. Boundaries of the three geographical regions are indicated by thicker lines.

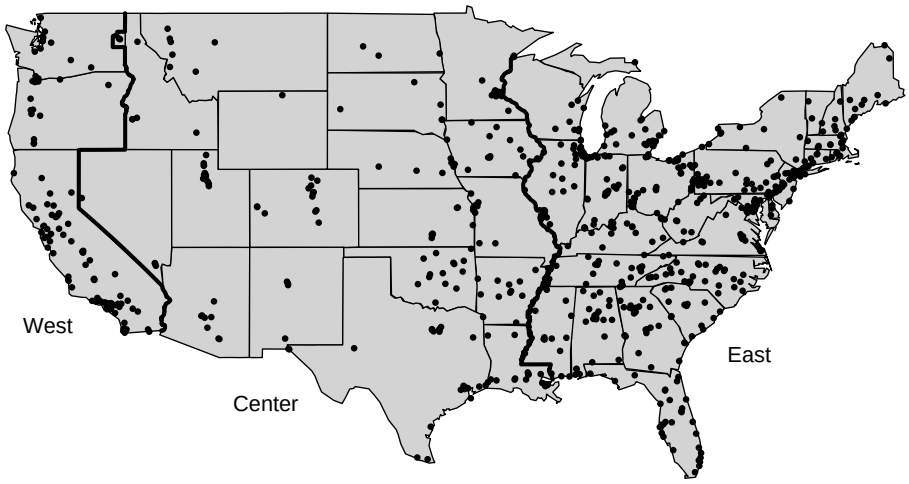

Figure 2: Average $PM_{2.5}$ concentrations over the previous year for the months from December 2000 to September 2006. Depicted are both the average across 540 monitors with complete time series in the continental U.S. (dotted-dashed line), as well as the average by region: the West (solid line), the center (dashed line), and the East (dotted line).

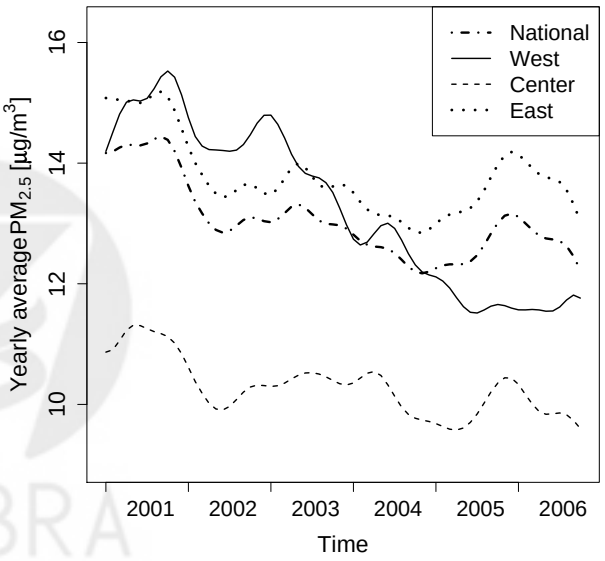

Figure 3: Monthly age-standardized mortality rates among Medicare enrollees from each of the three regions for December 2000 to September 2006. Mortality rates are standardized to the cohort study population in October 2003, the middle of the study period.

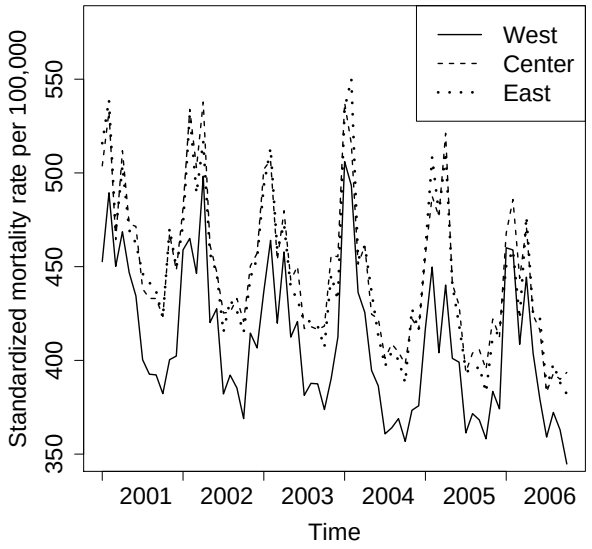

Figure 4: Estimated relative risk of dying in a given month per 10 $\mu g/m^3$ increase in average $PM_{2.5}$ concentrations during the previous year. Relative risk (RR) estimates based on the local coefficient β_1, the global coefficient β_2, and the overall coefficient β are shown with 95% confidence intervals.