

Air Pollution and Mortality: Estimating Regional and National Dose–Response Relationships

Francesca DOMINICI, Michael DANIELS, Scott L. ZEGER, and Jonathan M. SAMET

We analyzed a national data base of air pollution and mortality for the 88 largest U.S. cities for the period 1987–1994, to estimate relative rates of mortality associated with airborne particulate matter smaller than 10 microns (PM_{10}) and the form of the relationship between PM_{10} concentration and mortality. To estimate city-specific relative rates of mortality associated with PM_{10} , we built log-linear models that included nonparametric adjustments for weather variables and longer term trends. To estimate PM_{10} mortality dose–response curves, we modeled the logarithm of the expected value of daily mortality as a function of PM_{10} using natural cubic splines with unknown numbers and locations of knots. We also developed spatial models to investigate the heterogeneity of relative mortality rates and of the shapes of PM_{10} mortality dose–response curves across cities and geographical regions. To determine whether variability in effect estimates can be explained by city-specific factors, we explored the dependence of relative mortality rates on mean pollution levels, demographic variables, reliability of the pollution data, and specific constituents of particulate matter. We implemented estimation with simulation-based methods, including data augmentation to impute the missing data of the city-specific covariates and the reversible jump Markov chain Monte Carlo (RJMCMC) to sample from the posterior distribution of the parameters in the hierarchical spline model. We found that previous-day PM_{10} concentrations were positively associated with total mortality in most of the locations, with a .5% increment for a $10 \mu\text{g}/\text{m}^3$ increase in PM_{10} . The effect was strongest in the Northeast region, where the increase in the death rate was twice as high as the average for the other cities. Overall, we found that the pooled concentration–response relationship for the nation was linear.

KEY WORDS: Air pollution; Data augmentation; Generalized additive model; Hierarchical model; Natural cubic spline; Particulate matter; Relative rate.

1. INTRODUCTION

Epidemiologic time series studies conducted in cities around the world have consistently found associations between daily levels of airborne particulate matter smaller than 10 microns (PM_{10}), and daily numbers of deaths. These findings have raised concern about the public health effects of PM pollution (Schwartz 1994; Pope, Dockery, and Schwartz 1995a; American Thoracic Society and Bascom 1996a), and motivated reassessment of air quality standards in many countries, including the United States, the United Kingdom, and the European Union members. However, one key limitation of these studies has been use of data from a single or a few, possibly nonrepresentative, locations. The National Morbidity, Mortality, and Air Pollution Study (NMMAPS) addresses this limitation by assembling and analyzing a national data base that includes information on mortality, weather, and air pollution for the 88 largest metropolitan areas in the United States. The statistical framework estimates associations between air pollution and mortality (and morbidity) for the entire United States, within large regions and for particular cities (Samet,

Zeger, Dominici, Dockery, and Schwartz 1999; Samet et al. 2000b).

Initially, we analyzed data for the 20 largest U.S. cities, using a two-stage linear regression model for combining evidence from multiple locations (Dominici, Samet, and Zeger 2000). We then extended this analysis to estimate the shape of the PM_{10} mortality dose–response curve (Daniels, Dominici, Samet, and Zeger 2000). The 20-city analyses were constrained in model development and in their substantive findings by the relatively small number of cities. Further investigation is needed into the heterogeneity of the dose–response relationship of air pollution and mortality across cities and regions and on the modification of the effects of PM_{10} by factors such as copollutants, PM composition, city or regional measurement error, and population characteristics and susceptibilities.

In this article we extend the NMMAPS data base and analyses to include the 88 largest U.S. cities. The objectives of this article are (1) to combine information across these 88 locations to estimate regional and national relative rates of mortality from exposure to PM_{10} , (2) to explore heterogeneity of effects across broad geographic regions and determinants of heterogeneity, and (3) to estimate regional and national air pollution mortality dose–response curves. To determine whether city-specific factors can explain variability in the relative rates of mortality for PM_{10} , we have collected data on a set of city-specific variables: demographic characteristics, copollutant levels, precision of the air pollution measurements, and particle size distribution. A subset of these city-specific variables are missing in some cities, and thus a strategy for imputing missing data is needed. To address objectives (1) and (2), we develop a three-stage linear regression model with data augmentation to handle missing data in the city-specific

Francesca Dominici is Assistant Professor, Department of Biostatistics, Bloomberg School of Public Health, The Johns Hopkins University, Baltimore, MD 21205 (E-mail: fdominic@jhsph.edu). Michael Daniels is Assistant Professor, Department of Statistics, Iowa State University, Ames, IA (E-mail: mdaniels@iastate.edu). Scott L. Zeger is Professor and Chairman, Department of Biostatistics, Johns Hopkins University, Baltimore, MD 21205 (E-mail: szeger@jhsph.edu). Jonathan M. Samet is Professor and Chairman, Department of Epidemiology, Johns Hopkins University, Baltimore, MD 21205 (E-mail: jsamet@jhsph.edu). The research described in this article was partially supported by a contract and grant from the Health Effects Institute (HEI), an organization jointly funded by the Environmental Protection Agency (EPA; R824835) and automotive manufacturers. The contents of this article do not necessarily reflect the views and policies of HEI, EPA, or motor vehicle or engine manufacturers. Funding for Francesca Dominici was provided by the HEIs Walter A. Rosenblith New Investigator Award. Funding of Michael Daniels was provided by National Science Foundation grant DMS9816630. Funding was also provided by Johns Hopkins Center in Urban Environmental Health grant 5P30ES03819-12. The authors thank John Bachmann of the EPA for kindly providing us with the data on particulate matter composition, Giovanni Parmigiani for comments and suggestions on the statistical models, and Ivan Coursac for assistance with data base development.

covariates. To address (3), we use a two-stage spline model to estimate the shape of the PM_{10} mortality dose–response curves within each region.

Section 2 describes the data base of air pollution, mortality, and meteorological data from 1987 to 1994 for the 88 U.S. cities used in this analysis. Section 3 introduces the hierarchical regression model with data augmentation for combining information on the PM_{10} mortality associations across cities within regions and across regions. Section 4 introduces the hierarchical spline model used to estimate regional and national PM_{10} mortality curves. Section 5 summarizes the findings and presents results of model comparisons, model checking, and sensitivity to prior distributions. Finally, Section 6 discusses our findings. Details on the implementation of the reversible jump Morkov Chain Monte Carlo (RJCMCMC) for model fitting are presented in the Appendix.

2. DATA

Figure 1 shows the locations of the selected 88 cities and the 7 geographical regions of interest. We obtained data on mortality, weather, and air pollution these areas from publicly available data sources. Daily mortality counts were obtained from the National Center for Health Statistics. Hourly temperature and dew point data were available from the National Climatic Data Center, as assembled in the EarthInfo CD data base (<http://www.sni.net/earthinfo/cdroms/>). The daily time series of PM_{10} data for each city were obtained from the Aerometric Information Retrieval Service (AIRS) data base maintained by the U.S. Environmental Protection Agency (<http://www.wpa.gov/airs/airs.html>). One-year (1999) average $PM_{2.5}$ concentrations for each city were provided by Dr. John Bachmann of the Environmental Protection Agency. We also

collected information on multiple city-specific factors from the 1990 CensusCD (email: info@censuscd.com).

For the spatial analysis, we grouped the 88 counties into seven geographic regions (Northwest, Upper Midwest, Industrial Midwest, Northeast, Southern California, Southwest, Southeast), following the stratification used in the 1996 Review of the National Ambient Air Quality Standard for Particulate Matter (Environmental Protection Agency 1996). A detailed description of this database is given elsewhere (Samet et al. 1999; Samet et al. 2000b).

3. HIERARCHICAL LINEAR MODEL WITH DATA AUGMENTATION

In this section we introduce a three-stage hierarchical model with data augmentation to estimate regional and national relative rates of mortality for particulate matter, taking into account the missing data in the city-specific predictors.

In the first stage of the model, we fit a log-linear generalized additive model (Hastie and Tibshirani 1990), where the outcome variable, Y_{ct}^r , is the total number of deaths on a particular day t , in city c , within region r and the exposure variable, PM_{ct-1} , is the previous day's PM_{10} level:

$$Y_{ct}^r | \mu_{ct}^r \sim \text{Poisson}(\mu_{ct}^r), c = 1, \dots, C^r, \\ r = 1, \dots, R, t = 1, \dots, T, \quad (1)$$

$$\log \mu_{ct}^r = \beta_c^r PM_{ct-1} + \eta_c X_t.$$

Here C^r , R , and T denote the number of cities within each region r , the number of regions, and the number of days;




Figure 1. Map of the 88 Largest U.S. Cities With the 7 Geographical Regions of Interest. The gray scales are proportional to the estimated log relative rates of mortality, which show a -4% to 4% increase in mortality per $10 \mu g/m^3$ increase in PM_{10} . The circles' areas are proportional to the statistical precisions of the estimates. The larger circles show less statistical uncertainty. The circles with the black outline denote the relative rates that are statistically significant.

$\mu_{ct}^r = E[Y_{ct}^r]$; X_t is the t th row of the design matrix for the confounding factors (e.g., long-term trends and seasonality in the mortality time series, and weather variables); and $\boldsymbol{\eta}_c$ is the corresponding vector of coefficients. The potential confounding variables and the rationale for their inclusion, are listed in Table 1. Justification for selecting the smooth functions to control for longer-term trends, seasonality, and weather and sensitivity analyses with respect to the lag structure of the exposure variable have been given by Samet, Zeger, and Berhane (1995), Samet, Zeger, Kelsall, Xu, and Kalkstein (1997), Kelsall, Samet, and Zeger (1997), Samet, Dominici, Curriero, Coursac, and Zeger (2000a), and Dominici et al. (2000).

In several locations, a high percentage of days had missing PM_{10} values for which measurements are generally required only every six days. Because there are missing values of some predictor variables are missing on some days, we restricted analyses to days with no missing values across the full set of predictors.

At the second stage, we describe the heterogeneity of the city-specific effects within regions, assuming that

$$\boldsymbol{\beta}^r \mid \alpha_0^r, \boldsymbol{\alpha}, \mathbf{Z}^r, \sigma^2 \sim N_{C^r}(\alpha_0^r \mathbf{j}^r + \mathbf{Z}^r \boldsymbol{\alpha}, \sigma^2 I), \quad r = 1, \dots, R, \quad (2)$$

where $\boldsymbol{\beta}^r = [\beta_1^r, \dots, \beta_{C^r}^r]$ is the collection of true PM_{10} coefficients for the C^r cities in region r , \mathbf{Z}^r is a matrix of dimension $C^r \times p$ with the c th row $[Z_{c1}^r, \dots, Z_{cp}^r]$ denoting the centered covariates for the cities belonging to region r , α_0^r is the regional air pollution effect when all of the covariates are centered at their mean values, \mathbf{j}^r is a vector of length C^r having all elements equal to 1, and $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_p]$ is the vector of the second-stage regression coefficients (i.e., α_j measures the change in β_c^r per unit of change in the city-specific covariate Z_{cj}^r), and σ^2 measures the variance of the β_c^r 's within each region. The second-stage covariates are included in the design matrix \mathbf{Z} , and the rationale for their inclusion are summarized in Table 2. Details on the exploratory analyses that led to the selection of these variables appear elsewhere (Samet et al. 2000b).

At the third stage of the model, we investigated heterogeneity of the regional air pollution effects (α_0^r) across regions; we assume

$$\alpha_0^r \mid \alpha_0, \tau^2 \sim N(\alpha_0, \tau^2). \quad (3)$$

Here α_0 is the overall relative rate of mortality for PM_{10} , and τ^2 measures the variance of α_0^r across regions.

Because the vector $\boldsymbol{\eta}_c$ corresponding to the factors listed in Table 1 is highly dimensional (its dimension is 118), the computational demand of a full Bayesian approach—that is, simulating from the joint posterior distributions of $\boldsymbol{\beta}_c^r$ and $\boldsymbol{\eta}_c$ and then integrating over the $\boldsymbol{\eta}_c$ to obtain the marginal posterior distributions of the β_c^r —is extremely laborious. The computation becomes even more intensive when we implement data augmentation to impute the missing city-specific covariates. Therefore, let $\hat{\boldsymbol{\beta}}^r = [\hat{\beta}_1^r, \dots, \hat{\beta}_{C^r}^r]$ and $V^r = \text{diag}(v_1^r, \dots, v_{C^r}^r)$ be the maximum likelihood estimates (MLEs); their sampling variances are obtained by fitting the city-specific model (1) for city c and region r . We simplify the computation substantially by replacing the first stage of the model with the MLE-based normal approximation to the likelihood function,

$$\hat{\boldsymbol{\beta}}^r \sim N_{C^r}(\boldsymbol{\beta}^r, V^r). \quad (4)$$

Because of the large number of days with air pollution and mortality measurements within each city, we found that the MLE-based normal approximation to the likelihood is adequate, as discussed in Section 5.1.

This analysis is complicated by missing data for the second-stage variables in a subset of cities. (The percentages of cities with missing data for each variable are listed in Table 2.) We impute the missing data by implementing data augmentation (Tanner 1991) within the Gibbs sampler (Gelfand and Smith 1990), which requires specification of the conditional distributions of vectors of covariates that are missing (Z_{cm}^r) given the vector of covariates that are observed (Z_{co}^r). We assume that

$$Z_{\text{cm}}^r \mid Z_{\text{co}}^r, \boldsymbol{\theta}, \boldsymbol{\Sigma} \sim N_{l_c}(\boldsymbol{\mu}_o, \boldsymbol{\Sigma}_{\text{mm.o}}), \quad (5)$$

where

$$\begin{aligned} \boldsymbol{\mu}_o &= \boldsymbol{\theta}_m + \boldsymbol{\Sigma}_{\text{mo}} \boldsymbol{\Sigma}_{\text{oo}}^{-1} (Z_{\text{co}}^r - \boldsymbol{\theta}_o), \\ \boldsymbol{\Sigma}_{\text{mm.o}} &= \boldsymbol{\Sigma}_{\text{mm}} - \boldsymbol{\Sigma}_{\text{mo}} \boldsymbol{\Sigma}_{\text{oo}}^{-1} \boldsymbol{\Sigma}_{\text{om}} \end{aligned}$$

and $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are the mean and covariance matrix of the p vector $Z_c^r = [Z_{\text{cm}}^r, Z_{\text{co}}^r]$.

The model specification is completed with the selection of prior distributions for the parameters at the top level of the hierarchy. We assume a priori that these parameters are independent, and, except for the within-region and between-region variance components (σ^2 and τ^2), we choose vague conjugate priors with large variances. For σ^2 and τ^2 , we assume a half-normal prior distribution, which gives moderate weight to complete homogeneity while also allowing for the possibility of more substantial heterogeneity (Pauler and Wakefield

Table 1. Potential Confounding Factors in the Estimation of the City-Specific Relative Rates Associated With Particulate Air Pollution Levels, and the Rationale for Their Inclusion in the Model

Predictors	Primary reasons for inclusion
Indicator variables for the three age groups	To allow for different baseline mortality rates within each age group
Indicator variables for the day of the week	To allow for different baseline mortality rates within each day of the week
Smooth functions of time with 7 degrees of freedom (df)/yr	To adjust for long-term trends and seasonality
Smooth functions of temperature with 6 df	To control for the known effects of weather on mortality
Smooth functions of dewpoint with 3 df	To control for the known effects of humidity on mortality
Separate smooth functions of time (2 df/yr) for each age group contrast	To separately adjust for seasonality within each age group

Table 2. Second-Stage Variables and the Rationale for Their Inclusion in the Model

Predictors	Primary reasons for inclusion	Percent missing ^a
$\log \overline{PM}_{10}$	Possibility of a saturation effect	0
$\log \overline{O}_3$	To explore modification of the effect of one pollutant to another	11
$\log \overline{NO}_2$	To explore modification of the effect of one pollutant to another	35
$\log \% \text{NoHS}$	Potential heterogeneity of the effects associated with socio-demographic factors	0
$\text{logit}MCC$	Potential heterogeneity of the effects associated with the varying quality of the exposure	21
$PM_{2.5}/PM_{10}$	Hypothesized health effects of fine particles	0

NOTE: \overline{PM}_{10} , \overline{O}_3 , and \overline{NO}_2 denote the mean level of PM_{10} , mean level of ozone (O_3), and mean level of nitrogen dioxide (NO_2) over the period 1987–1994; %NoHS is percentage of persons lacking a high school degree; and MCC is median of all pairwise correlations of the PM_{10} measurements among the location-specific monitors.

^a The percentage of cities with missing data.

2000). In Section 5.1 we explore the consequences of this assumption using a sensitivity analysis of the posterior results to the prior specification. To obtain posterior distributions of all of the parameters of interest, we implement a Gibbs sampler with data augmentation to generate values from the posterior distribution of the unknown parameters and missing predictors.

4. HIERARCHICAL SPLINE MODEL

In this section we introduce a hierarchical spline model to estimate regional PM_{10} mortality dose–response curves. We allow for additional flexibility in the pollution–mortality relationship by replacing the linear term $\beta'_c PM_{ct-1}$ in the city-specific model (1) with $S(PM_{ct-1}, \text{knots})$ where S is a natural cubic spline with an unknown number of knots, k , at unknown locations $\nu = (\nu_1, \dots, \nu_k)$, and with boundary knots fixed at 0 and $100 \mu\text{g}/\text{m}^3$. We allow for up to two knots ($k \leq 2$) at possible locations $\{20, 25, \dots, 70\} \mu\text{g}/\text{m}^3$, and we assume that the number and location of knots is the same across cities within a region but may vary from region to region. We base this assumption on the increasing regional homogeneity of particulate pollution, reflecting the regional impact of such sources as power plants or mobile sources in large metropolitan areas. The seven regions designated by the Environmental Protection Agency reflect the increasing homogeneity within their bounds (Environmental Protection Agency 1996).

We assume

$$\begin{aligned} \log \mu'_{ct} &= S(PM_{ct-1}, \text{knots}) + \eta_c X_t \\ &= \phi'_c B_t + \eta_c X_t, \end{aligned} \quad (6)$$

where conditionally on k and ν , B_t is the t row of the design matrix for the natural cubic splines of the PM_{10} variable, ϕ'_c is the corresponding vector of coefficients, and X_t and η^c are the same as in model (1). Let $\hat{\phi}'_c$ and W'_c be the MLE of the vector of coefficients corresponding to the splines and their sampling covariance matrix for city c in region r obtained by fitting model (6). As in the hierarchical linear model of

Section 3, we replace the first stage of the model with a MLE-based normal approximation of the likelihood function,

$$\hat{\phi}'_c | \phi'_c \sim N_k(\phi'_c, W'_c), \quad c = 1, \dots, C^r. \quad (7)$$

To describe between-city variation of the dose–response curves within regions, we add a second level to the model,

$$\phi'_c | \delta^r, D^r \sim N_k(\delta^r, D^r), \quad r = 1, \dots, R, \quad (8)$$

where δ^r denotes the vector of coefficients of the regional dose–response curve and D^r denotes the between-city within-region covariance matrix.

We did not include city-specific covariates at the second level for two reasons. First, this would substantially increase the number of parameters in the model. Second, we expect that most of the heterogeneity would be explained by the nonlinearity allowed in the spline models, and that the inclusion of city-specific covariates would not contribute in the reduction of the between-city variability.

Because of the computational burden, we estimated dose–response curves separately within each region. A national dose–response curve is estimated by directly pooling information across the 88 cities, assuming that there is little between-region variability.

With regard to the prior specification, we assumed that the number of knots k and their locations ν have discrete uniform priors. Similarly for δ^r , we assumed a flat prior. For the prior on D^r , a computationally convenient choice is to place a constrained Wishart density on B'_0 , where $B'_0 = W_0^{1/2}(W_0 + D^r)^{-1}W_0^{1/2}$ and $W_0 = (1/C^r) \sum_{c=1}^{C^r} W'_c$ (Everson and Morris 2000).

To sample from the posterior distribution of the number of knots, their locations, and all other unknown parameters, we implemented an RJMCMC algorithm. Details on the implementation of the RJMCMC are provided in the Appendix. We assessed convergence by diagnostic methods (Raftery and Lewis 1992; Gelman and Rubin 1992; Gilks, Richardson, and Spiegelhalter 1996) in CODA (Best, Cowles, and Vines 1995).

5. RESULTS

Figure 1 shows the locations of the 88 cities, the magnitude of the estimated relative rates of mortality (color scale), and their statistical precisions (areas of the circles). We found that previous-day PM_{10} concentrations were positively associated with total mortality counts for most of the locations.

Figure 2 shows MLE and 95% confidence intervals of the log-relative rates of mortality per $10 \mu\text{g}/\text{m}^3$ increase in PM_{10} for each location. These estimates were obtained by fitting the semiparametric log-linear model (1) to the data for each location independently. The estimates of the city-specific effects ranged from highs of 2% per $10 \mu\text{g}/\text{m}^3$ in Oakland and 1% per $10 \mu\text{g}/\text{m}^3$ in New York City and San Diego to negative estimates in 22 of the 88 cities, including Dallas-Fort Worth, Little Rock, and Modesto.

The solid squares and the solid circles with the bold segments denote the posterior means and 95% posterior intervals of the pooled regional effects with and without covariate adjustment. Adjusted and unadjusted regional effects were




Figure 2. MLEs and 95% Confidence Intervals of the Log-Relative Rates of Mortality per 10 $\mu\text{g}/\text{m}^3$ Increase in PM_{10} for Each Location. The solid and the square circles with the bold segments denote the posterior means and 95% posterior intervals of the pooled regional effects without and with covariate adjustments. At the bottom, marked with triangles and bold segments, are the overall effects for PM_{10} for all the cities without and with covariate adjustments.

similar in most regions. The pooled regional estimates of the PM_{10} effects varied somewhat across the regions and were estimated to be greatest in the Northeast, with a relative rate of .9% per 10 $\mu\text{g}/\text{m}^3$ (95% CI .58, 1.31). On the far right, marked with triangles and bold segments, are the overall effect estimates for PM_{10} for all of the cities with and without covariate adjustment. The overall effect with covariate adjustment was slightly larger than the overall effect without covariate adjustment (without adjustment, a posterior mean = .43% increase

in mortality per 10 $\mu\text{g}/\text{m}^3$ increase in PM_{10} , 95% posterior interval .06, .77; with adjustment, a posterior mean = .55% increase in mortality per 10 $\mu\text{g}/\text{m}^3$ increase in PM_{10} , 95% posterior interval .10, .98).

Figure 3 shows boxplots of the posterior distributions of the second-stage regression coefficients α . We found some evidence for modification of the log-relative rate of mortality associated with PM_{10} levels by $\log \overline{\text{PM}}_{10}$ level in a direction implying greater shorter-term (acute) effects at lower




Figure 3. Boxplots of the Posterior Distributions of the Second-Stage Regression Coefficients.

longer-term (chronic) levels of PM_{10} . The negative association between the relative rate of mortality and $\log \overline{PM}_{10}$ might indicate a saturation effect of air pollution on mortality. One explanation of this phenomenon is that the pool of susceptible individuals could be depleted in cities with higher average PM concentrations. If this were the case, then the short-term effects of air pollution on mortality would be lower than in cities with a lower PM average and a relatively larger pool of susceptible individuals.

The positive coefficients for O_3 and NO_2 levels (although with a large posterior standard error) might indicate a potential confounding or modifying effect, or both, of the gaseous pollutants on the relative rate of mortality. The positive coefficient for %NoHS (although with a large posterior standard error) might imply that cities with lower socioeconomic status have a higher PM_{10} effect on mortality. The negative coefficient for $PM_{2.5}/PM_{10}$ is puzzling and inconsistent with previous studies, which found that fine particles, $PM_{2.5}$, are more likely than coarse particles to be responsible for health effects ($PM_{10} - PM_{2.5}$) (Levy, Hammitt, and Spengler 2000). However, we found when the $PM_{2.5}/PM_{10}$ ratio is the only predictor included in the second-stage regression, the estimated regression coefficient for $PM_{2.5}/PM_{10}$ increased from $-.7$ (std = .91) in the multivariable model, as shown in Figure 3, to $-.057$ (std = .58). In summary, except for $\log \overline{PM}_{10}$, we did not find strong evidence of effect modification by the other predictors included in the model.

Posterior distributions of the within- and between-region variances on the \log_{10} scale ($\log_{10} \sigma^2$ and $\log_{10} \tau^2$) are shown in Figure 4. The second row of numbers represents the original scale of σ^2 and τ^2 . A posteriori, the within-region variance was much smaller than the between-region variance, indicating that the effects of PM_{10} on mortality tend to be similar for cities within the same geographical region and relatively more heterogeneous for cities in two different regions. However, because of the small number of regions, the posterior distribution of τ^2 tended to be quite sensitive to the prior specification. Section 5.1 explores the sensitivity of these findings to the prior assumption with regard to heterogeneity across cities and across regions.




Figure 4. Posterior and Prior Distributions of the Within- and Between-Region Variances on the Log Scale ($\log_{10} \sigma^2$ and $\log_{10} \tau^2$). The bottom row of numbers represents the original scale of σ^2 and τ^2 . (a) The posterior distribution of $\log_{10} \sigma^2$ under the half-normal prior. The curve with the bold dots represents the half-normal prior in \log_{10} scale. (b) The posterior distribution of $\log_{10} \tau^2$ under the half-normal prior. The curve with the bold dots represents the half-normal prior in \log_{10} scale. The dotted dark line represents the posterior distribution of $\log_{10} \tau^2$ under an inverse gamma prior $IG(3, 1)$, and the curve with the empty dots represents the inverse gamma prior $IG(3, 1)$ in \log_{10} scale.

Regional PM_{10} mortality dose–response curves are shown in Figure 5 with bold solid lines. These curves were estimated separately within each region, allowing for an unknown number and location of knots. The curves with empty dots were estimated by using a simplified version of the RJMCMC algorithm with one fixed knot at $40 \mu\text{g}/\text{m}^3$. At the bottom right are the national dose–response curves. The bold line represents the estimated national dose–response curve, allowing for unknown numbers and locations of knots and obtained by directly pooling information across the 88 cities. This approach assumes little between-region variability, which is not reflected in the hierarchical linear model. We also estimated a national dose–response curve by pooling the fixed knot regional curves with a two-stage Bayesian hierarchical model (Daniels et al. 2000). This national curve, represented by the empty dots, overlaps with the national curve estimated using RJMCMC. To allow for a direct comparison with the spline model, the regional linear curves are obtained under the hierarchical linear model without any shrinkage across regions.

The national dose–response curve, obtained by combining information across all the cities, is clearly linear. At the regional level, the data from cities in several regions (Northwest, Southwest, Upper Midwest, and Southeast) indicate some modest departure from a linear model. In particular, evidence in the Southwest and Southeast indicate a change point of about $30 \mu\text{g}/\text{m}^3$. The Northwest and Upper Midwest regions show a leveling off (saturation effect) at higher PM levels. This is consistent with the negative association between the relative rate of mortality and $\log \overline{\text{PM}}_{10}$ found in the hierarchical linear model. However, these dose–response curves have considerable uncertainty, and the pointwise confidence bands are consistent with a linear relationship.

5.1 Model Checking, Model Comparison, and Sensitivity Analysis

In this section we check modeling assumptions, discuss results of model comparisons, and investigate the impact of the prior distribution on our results.

5.1.1 Model Checking. In Section 3 we replaced the first stage of the city-specific model (1) with an MLE-based normal approximation to the likelihood function (4). To verify the adequacy of this approximation, we first selected the five cities with the smallest number of PM_{10} measurement days (sample sizes between 338 and 376). Within each city, we implemented a fully Bayesian analysis of the city-specific model (1) and estimated the joint posterior distribution of all unknown parameters (relative rates of mortality and coefficients of the splines) by importance sampling (Hammersley, Handscomb, and Muller 1966; Geweke 1989; Wakefield, Gelfand, and Smith 1991). We used proper but vague priors for all of the coefficients. Samples from the marginal posterior distribution $p(\beta_c^r | \text{data})$ and samples from the normal approximation to the likelihood (4) were very similar (results not shown).

5.1.2 Model Comparison. To investigate the impact of the data augmentation approach on our results, we compared Bayesian and frequentist estimates of α . The first was

obtained under the hierarchical model with data augmentation as described in Section 3, and the second under the linear mixed-effects model with EM algorithm applied only to the cities with complete data (48 total). More specifically, we compared estimates of α under the following three scenarios for departure from the hierarchical linear model with data augmentation:

- NR: Fit a linear mixed-effects model by the Newton–Raphson (NR) algorithm (Lindstrom and Bates 1988) to the complete data. Here $\hat{\beta}^r$ is the outcome; α_0^r , $r = 1, \dots, R$ are the random effects; and α are the fixed effects.
- BC1: Fit the hierarchical linear model to the complete data, ignoring the statistical error in the estimated $\hat{\beta}^r$ s, that is, assuming $\beta^r = \hat{\beta}^r$.
- BC: Fit the hierarchical linear model to the cities with complete data.

Point estimates and 95% confidence intervals of the second-stage regression coefficients α under each scenario are shown in Figure 6. Results under the baseline model had narrower confidence intervals, indicating that parameters can be estimated more precisely by using all of the 88-cities data with data augmentation than by using only the complete data. This is particularly true for the $\log \overline{\text{PM}}_{10}$ coefficient. Here, the association between the relative rate of mortality and the $\log \overline{\text{PM}}_{10}$ is negative when we include all of the cities than when we consider only the cities with complete data, because cities with missing data in the other covariates tend to have lower $\log \overline{\text{PM}}_{10}$ and larger relative mortality rates than cities with complete data.

The hierarchical spline and linear models can be compared easily, because the linear model is a special case of the spline model with 0 knots. We did so by inspecting the posterior distribution of the number of knots k obtained under RJMCMC (see Table 3). At a national level, the data clearly supported the hypothesis of linearity with a posterior probability of $k = 0$ equal to 1. Most of the regions lent some support to the linear model. For example, the Industrial Midwest and Northeast had posterior probabilities of $k = 0$ equal to .94 and .44. Only two regions had probabilities of more than .10 for the two-knot model, with five regions lending the most support to the one-knot model. The linearity of the national dose–response curve most likely reflects averaging out of the nonlinearities at the regional and city levels when the data are pooled nationally. In addition, none of the confidence boundaries at the regional level was inconsistent with a linear relationship.

5.1.3 Sensitivity to the Priors for τ^2 and σ^2 . The posterior distributions for the regression parameters are dominated by the likelihood functions given the vague prior assumption. Inferences about the degree of heterogeneity in pollution effects among cities and regions however were sensitive to the prior assumptions about σ^2 and τ^2 .

Our strategy for investigating the impact of the prior distribution on our results was based on inspecting the posterior distributions of the unadjusted overall PM_{10} effects on total mortality α_0 and the between-region variance τ^2 under four scenarios for departure from the baseline prior distribution for the two variance components σ^2 and τ^2 .

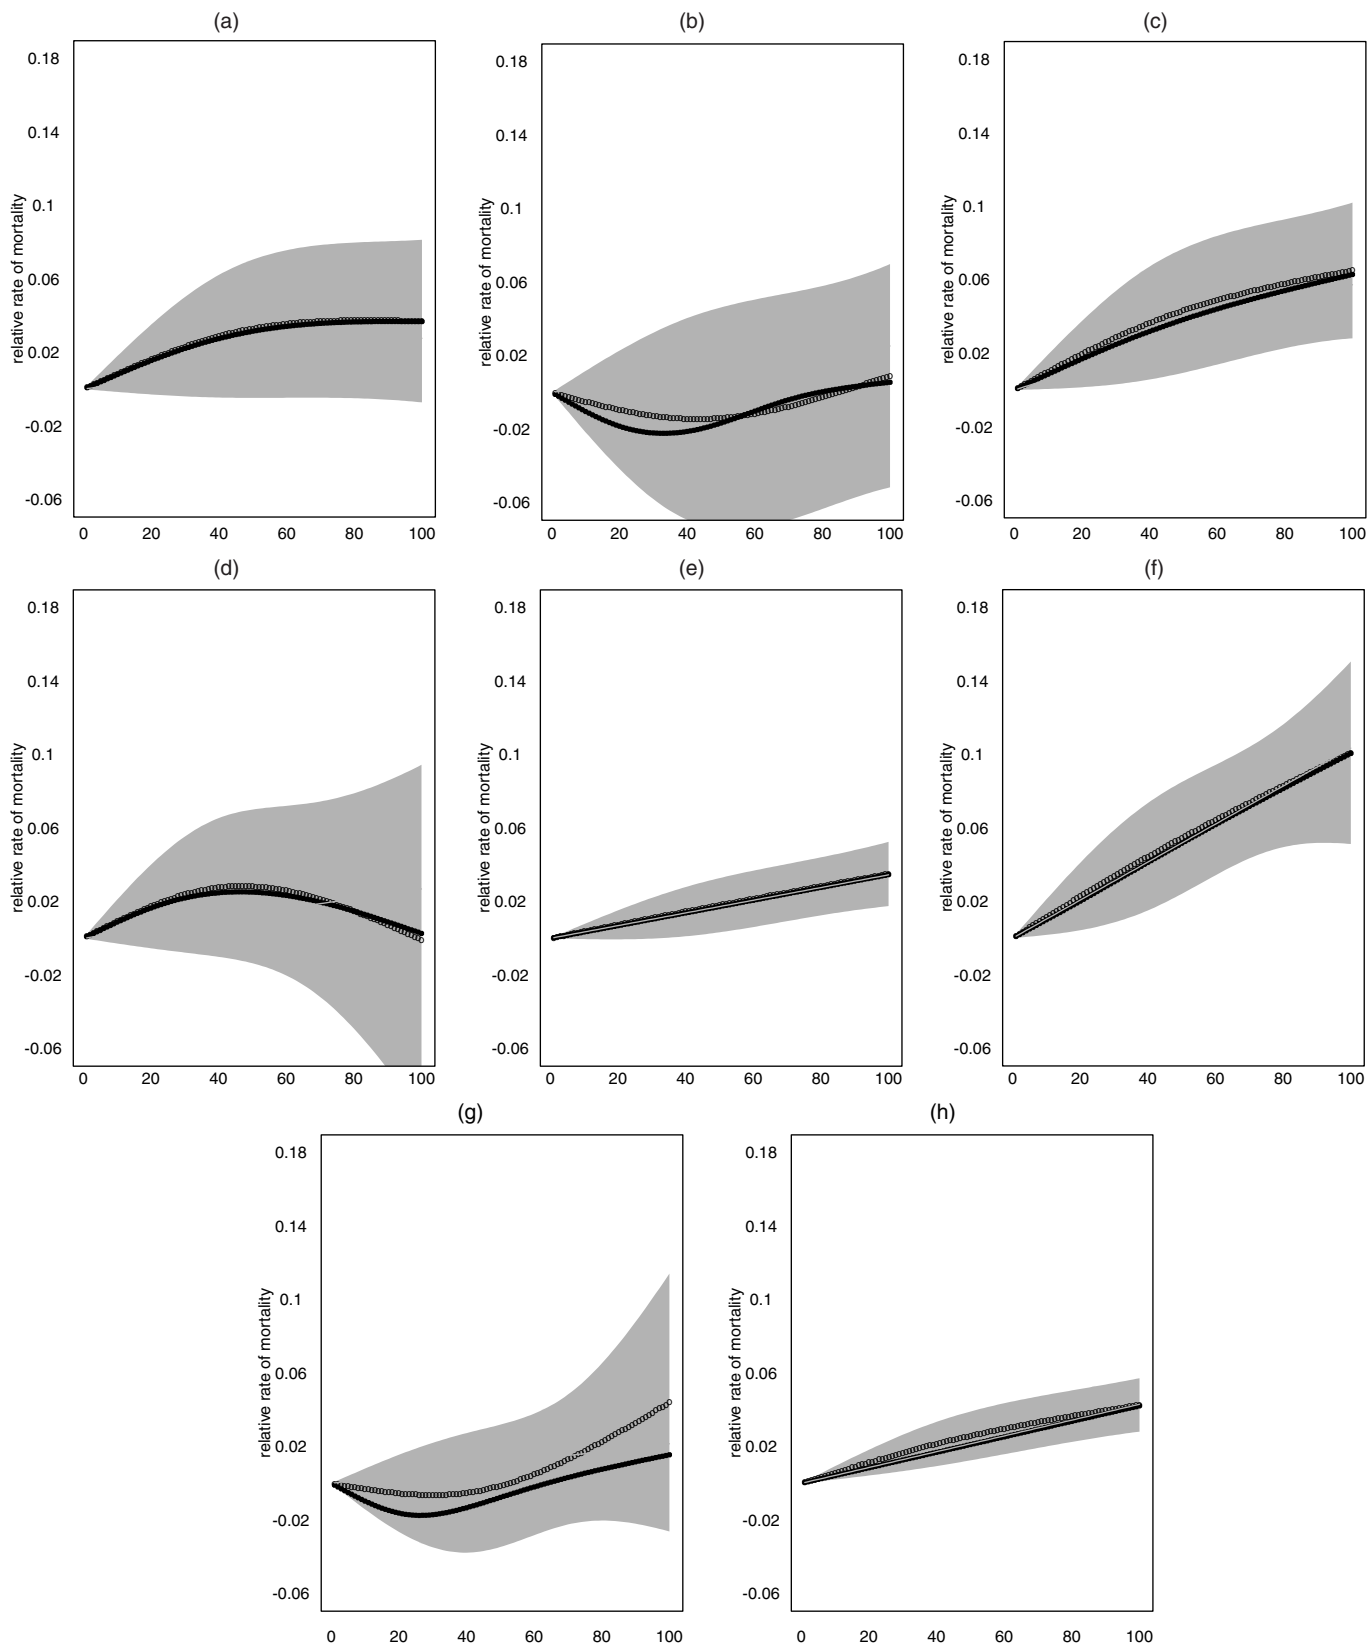


Figure 5. Regional and National PM_{10} Mortality Dose-Response Curves for (a) Northwest, (b) Southwest, (c) Southern California, (d) Upper Midwest, (e) Industrial Midwest, (f) Northeast, (g) Southeast, and (h) Overall. The x-axes are PM_{10} Levels in $\mu g/m^3$. The vertical scale can be interpreted as the relative rate of mortality as a function of PM_{10} . The solid black curves are obtained by fitting the spline model with the RJMCMC and allowing for an unknown number and location of knots. The curves with the empty dots are obtained by setting one knot at $40 \mu g/m^3$ and fitting the spline model conditional on the restricted MLE of the between-city covariance matrix. The linear curves are obtained by fitting the hierarchical linear model with a Gibbs sampler without borrowing strength across regions. The shaded area denotes the 95% confidence bands for the curve with a fixed knot.




Figure 6. Point Estimates and 95% Confidence Intervals of the Second-Stage Regression Parameters α Estimated With the Following Models: NR, mixed linear effect model for the complete case with NR algorithm and $\beta' = \hat{\beta}'$; BC1, hierarchical linear model for the complete case and $\beta' = \hat{\beta}'$; BC, hierarchical linear model for the complete case; B, hierarchical linear model with data augmentation (baseline). Horizontal lines are placed at 0.

The four scenarios and the posterior quantiles of α_0 under each prior specification, are summarized in Table 4. Posterior means and IQR of α_0 were not sensitive to the prior specifications of σ^2 and τ^2 . Prior and posterior distributions on τ^2 are shown in Figure 4. The curve with the bold dots represents the half-normal prior in \log_{10} scale, and the histogram represents the posterior distributions of $\log_{10} \tau^2$ under the half-normal prior. The curve with the empty dots represents the inverse γ prior $IG(3, 1)$ in \log_{10} scale, and the dotted dark line represents the posterior distribution of $\log_{10} \tau^2$ under the inverse γ prior. As expected, because of the small number of regions, the posterior distribution of τ^2 was sensitive to the prior. The inverse γ prior (curve with the empty dots) assumed more heterogeneity and was more concentrated than the half-normal prior. Consequently, a posteriori, we estimated

Table 3. Posterior Distributions of the Number of Knots, k , Within Each Region and Overall

Region	$k = 0$	$k = 1$	$k = 2$
Northwest	.10	.86	.04
Southwest	.02	.63	.35
Southcal	.31	.68	.01
Upper Midwest	.19	.75	.06
Industrial Midwest	.94	.06	0
Northeast	.44	.48	.08
Southeast	.22	.13	.66
Overall	1.00	0	0

a larger between-region variability. In summary, the inverse Γ prior on one or both variance components produced a posterior distribution for τ^2 with a larger posterior standard deviation ($E[\tau | \text{data}] = .38$ with half-normal prior and $E[\tau | \text{data}] = .49$ with inverse γ), and city-specific relative rate estimates, which drew more heavily on data from each city and less heavily on data from other cities, yielding larger and more conservative confidence bands on the overall relative rate.

6. DISCUSSION

In the United States and other developed countries, air quality standards for public health protection have their basis in the scientific evidence on risks to health. In the United States, for example, the Administrator of the Environmental Protection Agency is required by the Clean Air Act to set standards for the major (so-called “criteria”) air pollutants that protect the public health within an “adequate margin of safety.” To fulfill this charge, information is needed on the risks to health at different levels of the pollutant(s) and on the form of the concentration–risk relationship.

The specification of statistical models used to estimate associations between air pollution and health requires an extensive preprocessing of variables and a series of analytic choices. In particular, for the city-specific model these choices were based mainly on the specification of variables to adjust for temporal trends, weather, specification of the lag structure, selection of pollutants and monitors, and stratification by age. These analytic choices and the sensitivity of the results to them have been a major focus of our research (Samet et al. 1995, 1997, 1998, 1999, 2000b; Kelsall et al. 1997; Dominici et al. 2000) and were briefly discussed in this article. The selection of the variables included at the second stage of the hierarchical model was also the result of extensive exploratory analyses summarized elsewhere (Samet et al. 2000b). The groups of variables presented there and initially considered as poten-

Table 4. Posterior Quantiles of the Unadjusted Overall PM_{10} Effect on Total Mortality, α_0 , Under Four Specifications of Prior Distributions on σ^2 and τ^2

Prior for σ^2	Prior for τ^2	Scenario	5%	25%	50%	75%	95%
$\sigma^2 \sim N(0, 1)I_{\sigma^2 > 0}$	$\tau^2 \sim N(0, 1)I_{\tau^2 > 0}$	Baseline	.06	.33	.43	.53	.77
$\sigma^2 \sim IG(3, 3)$	$\tau^2 \sim N(0, 1)I_{\tau^2 > 0}$	1	.09	.34	.47	.58	.85
$\sigma^2 \sim N(0, 1)I_{\sigma^2 > 0}$	$\tau^2 \sim IG(3, 1)$	2	.02	.29	.42	.55	.78
$\sigma^2 \sim IG(3, 3)$	$\tau^2 \sim IG(3, 1)$	3	.00	.30	.45	.60	.90

tial determinants of the heterogeneity of the effects across cities were (1) mean levels of pollutants, temperature, and dew point; (2) total mortality rates; (3) sociodemographic variables; (4) variables related to urbanization; and (5) variables related to measurement error. Based on the correlation patterns between the estimates of the relative rates of mortality and the potential effect modifiers, we limited the number of variables in each of the five groups and selected the second-stage variables used in the article.

Our modeling approach uses the same set of confounding variables in the different locations, although the estimates of the confounding effects are city-specific. An alternative approach would have been to select these confounding variables differently within each city, conditionally on some desired optimality criteria. For example, we could have selected a city-specific number of degrees of freedom in the smooth functions of time, based on the inspection of the autocorrelation function of the residuals of the mortality time series. We decided against this approach for two reasons: (1) to ensure direct comparability of the city-specific coefficients at the second level of the hierarchical model and (2) to maintain a clear interpretation of the regional-specific relative mortality rates.

The strengths of these methods lie in the synthesis of evidence across broad regions and the characterization of heterogeneity of effect at the city or regional level. To guide policy development, concentration–response relationships need to be described with sufficient certainty (Barnett and O’Hagan 1997); this certainty can be gained by pooling the large amounts of publicly available data on mortality, air pollution, and potential confounders and modifiers. Our methods, although computationally intense, can be implemented on publicly available data. There are obvious extensions to other environmental health problems, such as morbidity from air pollution. The methods are weakened by the gaps in the publicly available data on air pollution, mortality, and city-specific characteristics and by the inherent limitations of these data. The air pollution data routinely available from required regulatory monitoring provide only the mass of particulate matter, with no information on chemical or other characteristics of the particles, such as size distribution. After the 1997 promulgation of new standards for $PM_{2.5}$, [monitoring was initiated for \$PM_{2.5}\$](#) , and, in some locations, particles are characterized by their chemical composition. Thus, future analyses of heterogeneity of effects may be more informative. Total mortality—the outcome measure used in these analyses—is a crude measure, undoubtedly including a substantial proportion of deaths not attributable to air pollution; however, findings for total mortality have obvious public health relevance. More specific categories of cause of death can be used—for example, cardiovascular and respiratory deaths. Although specificity is gained, precision is lower because of the smaller number of events. In addition, the cause of death information specified on the death certificate is not always reliable. The decennial censuses provide information on city-specific characteristics that have the advantage of ready availability, but provide only partial coverage of some potentially relevant characteristics, such as time-activity patterns.

Overall, we found that the pooled concentration–response relationship for the nation was approximately linear. There

was little evidence for deviation from linearity down to the lowest levels to support the existence of a threshold. This finding confirms the prior analysis by Daniels et al. (2000) of data from 20 U.S. cities (included within the 88 cities considered here). As recently summarized by Pope (2000), analyses of time series data in a number of locations, including London, Cincinnati, Birmingham, Utah Valley, and Shenyang, are also consistent with linearity, as are the findings of the two major long-term studies, the Harvard Six-Cities cohort and the American Cancer Society cohort (Dockery et al. 1993; Pope et al. 1995b). The consistent finding of a linear relationship, now confirmed at the national level, places a difficult burden on policy makers charged by the Clean Air Act with setting protective standards for public health that include a margin of safety.

We found indications of potential nonlinearity in four of the regions: the Northwest, Southwest, Upper Midwest, and Southeast. The analyses for the Northwest and Upper Midwest regions also suggest a leveling of the effect at higher PM values. This is consistent with the negative association between the city-specific relative rates of mortality and $\log \overline{PM}_{10}$ found in the hierarchical linear model. However, the uncertainty boundaries for these regions indicate compatibility of the data with a linear relationship, and we currently have no specific hypotheses explaining why these regions might have other than a linear dose–response curve.

In general, the city-specific characteristics did not explain much variation in the effect of PM across cities. However, we found evidence in the data suggesting that cities with lower mean PM_{10} levels tend to have larger relative rates of mortality from PM_{10} exposure. The negative association between the relative rate of mortality and the mean PM_{10} is consistent with our mortality displacement analyses (Zeger, Dominici, and Samet 1999; Dominici, McDermott, Zeger, and Samet 2001), which showed that the association between air pollution and mortality is larger at the longer-term variations in the time series, corresponding to a more chronic effect of air pollution on mortality, than at the short-term variations in the time series, corresponding to a more acute effect of air pollution on mortality. These results are also consistent with the findings of two long-term prospective cohort studies of air pollution and mortality, the Harvard Six-Cities Study (Dockery et al. 1993) and the American Cancer Society’s Cancer Prevention Study II (Pope et al. 1995b), which offer critical evidence of a long-term chronic effect of air pollution on mortality.

In summary, we offer a set of methods for combining evidence in publicly available databases on the concentration–response relationship for air pollution and mortality. A multi-stage analysis was used because a uniform analytic approach can be applied to each city, power is gained from the pooled sample, and heterogeneity can be explored across the locations.

Previous meta-analyses of the PM literature have provided pooled effects estimates, but have not addressed between-study variability that may be associated with analytical models, pollution patterns, and exposed populations (Schwartz 1994). The meta-analysis by Levy et al. (2000) was one of the first attempts to determine whether study-specific factors

can explain some of the variability of the relative rate of mortality associated with PM_{10} . However, meta-analyses of published studies of air pollution and health were limited by three factors: (1) the lack of a common analysis of the raw data, (2) the limited number of locations, and (3) the study selection bias. Our hierarchical approach overcomes these limitations (Dominici 2001).

In developing these methods, our intention was to provide evidence in a form that would be directly applicable to policy development, as shown by our finding that the concentration–response relationship for air pollution and mortality is linear with a high degree of certainty. With repeated application, our methods would offer an approach for tracking the health effects of air pollution over time as control measures are implemented; they should have application to other environmental health problems as well.

APPENDIX A: IMPLEMENTATION OF THE REVERSIBLE JUMP

We draw samples from the joint posterior distribution of (\mathbf{v}, k) by implementing the following RJMCMC algorithm:

1. Choose a type of move and do the move. That is, either add a knot, delete a knot, or move an existing knot. Implementation of this step follows that described by Denison, Mallick, and Smith (1998) and Di Matteo, Genovese, and Kass (2000).
2. Compute $\hat{\Phi}_c^r$ and W_c^r conditional on the number and location of knots. These estimated coefficient and sampling covariances change as the number and location of the knots change.
3. Compute \hat{D}^r , the restricted MLE, conditional on the number and location of knots.
4. Compute the acceptance probability to determine whether to accept (\mathbf{v}, k) for this iteration. The acceptance probabilities for the birth and deletion of a knot are given by α_b and α_d and are defined in the next section.
5. Repeat steps 1–4 M times.

For each iteration from the RJMCMC algorithm, we obtain a sample from the marginal posterior distribution of (\mathbf{v}, k) conditional on \hat{D}^r and on $\delta^{r(j)}$. We then sample δ from $[\delta^r | \hat{D}^{r(j)}, \mathbf{v}^{(j)}, k^{(j)}, \text{data}]$. This distribution is multivariate normal with covariance matrix $V = [\sum_c (W_c^r + \hat{D}^r)^{-1}]^{-1}$ and mean $V \times \sum_c (W_c^r + \hat{D}^r)^{-1} \hat{\Phi}_c^r$.

Note that we cannot simply average the coefficient of the spline curves across iterations as the number and location of the knots changes. We can obtain the posterior mean of the regional (or city-specific) curve by averaging the curves specified by δ^r ($\hat{\Phi}_c^r$) from each iteration. We do this by choosing a grid of PM_{10} values and then pointwise averaging the curves over the grid.

An alternative approach to fixing D^r at its restricted MLE is to sample B_0^r from $[B_0^r | \mathbf{v}, k, \text{data}]$, where $B_0^r = W_0^{1/2} (W_0 + D^r)^{-1} W_0^{1/2}$ and $W_0 = \frac{1}{C^r} \sum_{c=1}^{C^r} W_c^r$ (see Everson and Morris 2000 for details).

To show that the RJMCMC algorithm proposed earlier converges to the correct distribution, we need to show that the Markov chain satisfies detailed balance,

$$\pi(M_k)P(M_{k-1}|M_k) = \pi(M_{k-1})P(M_k|M_{k-1}). \quad (\text{A.1})$$

This proof follows very closely the proof of Di Matteo et al. (2000), and we use their notation. In (A.1), M_k denotes the parameters of the model with k knots and M_{k-1} denotes the parameters of the model with $k-1$ knots. $P(M_k|M_{k-1})$ is the transition probability for going from M_{k-1} to M_k and $\pi(M_k)$ is the posterior distribution of M_k ,

defined as

$$\pi(M_k) = \hat{L}(y|B_0)p(B_0|\nu, k)p(\mathbf{v}|k)p(k)/\hat{p}(y),$$

where \hat{L} is the approximation to the marginal likelihood for B_0 as defined by Daniels and Kass (1998), δ is integrated out, $\hat{p}(y)$ is the normalizing constant, $p(B_0, \nu, k) = p(B_0|\nu, k)p(\nu|k)p(k)$ is the joint prior distribution on the between-city covariance matrix, and the location and number of knots are as defined in Section 4.

Now define $b_k = P(k+1|k)$ to be the probability of adding a knot and $d_k = P(k-1|k)$ to be the probability of deleting a knot, when k knots are currently in the model. For our example, we set $b_k = d_k = .4$. Let α_d denote the acceptance probability for deleting a knot. Then

$$\begin{aligned} P(M_{k-1}|M_k) &= P(k-1|k)P(\text{delete } \nu_j|k)\hat{\pi}(B_0|\nu, k-1, \hat{\Phi})\alpha_d \\ &= d_k(1/k)\hat{\pi}(B_0|\nu, k, \hat{\Phi})\alpha_d, \end{aligned}$$

where

$$\alpha_d = \min\{1, A\}$$

with

$$A = \frac{\pi(M_{k-1})b_{k-1}(1/(n-(k-1)))}{\pi(M_k)} \frac{d_k 1/k}{\hat{\pi}(B_0|\nu, k, \hat{\Phi})} \frac{\hat{\pi}(B_0|\nu, k-1, \hat{\Phi})}{\hat{\pi}(B_0|\nu, k-1, \hat{\Phi})}.$$

Here $\hat{\pi}(B_0|\nu, k, \hat{\Phi})$ denotes the constrained Wishart density that most closely envelops the true full conditional distribution of B_0 (Everson and Morris 2000). To evaluate this constrained Wishart density, we need the normalizing constant (i.e., the probability that all of the eigenvalues of B_0 are less than or equal to 1). This can be computed by simulation (i.e., simulate from the same Wishart distribution, $\hat{\pi}$, without the constraint and then estimate the probability).

We now consider adding a knot, that is, going from a model with $k-1$ knots to a model with k knots. Here let α_b denote the acceptance probability for adding a knot,

$$\begin{aligned} P(M_k|M_{k-1}) &= P(k|k-1)P(\text{add } \nu_j|k-1)\hat{\pi}(B_0|\nu, k, \hat{\Phi})\alpha_b, \\ &= b_{k-1}(1/(n-(k-1)))\hat{\pi}(B_0|\nu, k, \hat{\Phi})\alpha_b, \end{aligned}$$

where

$$\alpha_b = \min\{1, B\}$$

with

$$B = \frac{\pi(M_k)}{\pi(M_{k-1})} \frac{d_k 1/k}{b_{k-1}(1/(n-(k-1)))} \frac{\hat{\pi}(B_0|\nu, k-1, \hat{\Phi})}{\hat{\pi}(B_0|\nu, k, \hat{\Phi})} = 1/A.$$

If $A < 1$, then $\alpha_d = A$ and $\alpha_b = 1$,

$$\begin{aligned} \pi(M_k)P(M_{k-1}|M_k) &= \pi(M_k)d_k(1/k)\hat{\pi}(B_0|\nu, k-1, \hat{\Phi})A \\ &= \pi(M_k)d_k(1/k)\hat{\pi}(B_0|\nu, k-1, \hat{\Phi}) \\ &\quad \times \frac{\pi(M_{k-1})b_{k-1}(1/(n-(k-1)))}{\pi(M_k)} \frac{\pi(B_0|\nu, k, \hat{\Phi})}{\hat{\pi}(B_0|\nu, k-1, \hat{\Phi})} \\ &= \pi(M_{k-1})b_{k-1}1/(n-(k-1))\hat{\pi}(B_0|\nu, k, \hat{\Phi}) \\ &= \pi(M_{k-1})P(M_k|M_{k-1}). \end{aligned}$$

Thus detailed balance is satisfied. Similar arguments can be used to show that moving a knot satisfies detailed balance.

If instead of sampling B_0 , we fix it at its restricted MLE, $\widehat{B}_0(\widehat{D})$, then the foregoing argument proceeds similarly. However, the product of the prior for B_0 and the likelihood approximation, $\widehat{L}(y|B_0)$, in $\pi(M_k)$ is now replaced by $\widehat{L}(y|\widehat{B}_0(\widehat{D}))$. In addition, the transition and acceptance probabilities remain the same, except for the exclusion of $\widehat{\pi}(B_0|\mathbf{v}, k, \widehat{\phi})$. The approach of fixing D at its restricted MLE estimate simplifies computations. However, it decreases the size of the penalty for models with more knots (see Di Matteo et al. 2000).

[Received December 2000. Revised August 2001.]

REFERENCES

- American Thoracic Society, and Bascom, R. (1996a), "Health Effects of Outdoor Air Pollution," Part 1, *American Journal of Respiratory and Critical Care Medicine*, 153, 3–50.
- (1996b), "Health Effects of Outdoor Air Pollution," Part 2, *American Journal of Respiratory and Critical Care Medicine*, 153, 477–498.
- Barnett, V., and O'Hagan, A. (1997), *Setting Environmental Standards: The Statistical Approach to Handling Uncertainty and Variation*, New York: Chapman and Hall.
- Best, N. G., Cowles, M. K., and Vines, K. (1995), "CODA: Convergence Diagnostics and Output Analysis Software for Gibbs Sampling Output," Technical Report, Version 0.30, University of Cambridge, MRC Biostatistics Unit.
- Daniels, M., Dominici, F., Samet, J. M., and Zeger, S. L. (2000), "Estimating PM10-Mortality Dose-Response Curves and Threshold Levels: An Analysis of Daily Time-Series for the 20 Largest U.S. Cities," *American Journal of Epidemiology*, 152, 397–412.
- Daniels, M., and Kass, R. (1998), "A Note on First-Stage Approximation in Two-Stage Hierarchical Models," *Sankhya*, Ser. B, 60, 19–30.
- Denison, D., Mallick, B., and Smith, A. (1998), "Automatic Bayesian Curve Fitting," *Journal of the Royal Statistical Society*, Ser. B, 60, 333–350.
- Di Matteo, I., Genovese, C., and Kass, R. (to appear), "Bayesian Curve Fitting With Free-Knot Splines," technical report, Carnegie Mellon University.
- Dockery, D., Pope, C. A., Xu, X., Spengler, J., Ware, J., Fay, M., Ferris, B., and Speizer, F. (1993), "An Association Between Air Pollution and Mortality in Six U.S. Cities," *New England Journal of Medicine*, 329, 1753–1759.
- Dominici, F. (in press), "Air Pollution and Health: What can we Learn from an Hierarchical Approach?" Invited Commentary, *American Journal of Epidemiology*.
- Dominici, F., McDermott, A., Zeger, S. L., and Samet, J. M. (2001), "Airborne Particulate Matter and Mortality: Time-scale Effects in Four U.S. Cities," *American Journal of Epidemiology*.
- Dominici, F., Samet, J. M., and Zeger, S. L. (2000), "Combining Evidence on Air Pollution and Daily Mortality From the Twenty Largest U.S. Cities: A Hierarchical Modeling Strategy" (with discussion), *Journal of the Royal Statistical Society*, Ser. A, 163, 263–302.
- Environmental Protection Agency (1996), "Review of the National Ambient Air Quality Standards for Particulate Matter: Policy Assessment of Scientific and Technical Information," OAQPS Staff Paper.
- Everson, P., and Morris, C. (2000), "Inference for Multivariate Normal Hierarchical Models," *Journal of the Royal Statistical Society*, Ser. B, 62, 399–412.
- Gelfand, A. E., and Smith, A. F. M. (1990), "Sampling-Based Approaches to Calculating Marginal Densities," *Journal of the American Statistical Association*, 85, 398–409.
- Gelman, A., and Rubin, D. B. (1992), "Inference From Iterative Simulation Using Multiple Sequences" (with discussion), *Statistical Science*, 7, 457–472.
- Geweke, J. (1989), "Bayesian Inference in Econometric Models Using Monte Carlo Integration," *Econometrica*, 57, 1317–1339.
- Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. (eds.) (1996), *Markov Chain Monte Carlo in Practice*, London: Chapman and Hall.
- Hammersley, J. M., Handscomb, D. C. A., and Muller, M. E. R. (1966), "Review of Monte Carlo Methods," *The Annals of Mathematical Statistics*, 37, 532–538.
- Hastie, T. J., and Tibshirani, R. J. (1990), *Generalized Additive Models*, New York: Chapman and Hall.
- Kelsall, J., Samet, J. M., and Zeger, S. L. (1997), "Air Pollution and Mortality in Philadelphia, 1974–1988," *American Journal of Epidemiology*, 146, 750–762.
- Levy, J., Hammitt, J., and Spengler, J. (2000), "Estimating the Mortality Impacts of Particulate Matter: What Can be Learned From Between-Study Variability?," *Environmental Health Perspectives*, 108.
- Lindstrom, M. J., and Bates, D. M. (1988), "Newton–Raphson and EM Algorithms for Linear Mixed-Effects Models for Repeated-Measures Data," *Journal of the American Statistical Association*, 83, 1014–1022.
- Pauler, D., and Wakefield, J. (2000), "Modeling and Implementation Issues in Bayesian Meta-Analysis," in *Meta-Analysis in Medicine and Health Policy*, pp. 231–254.
- Pope, C. A. (2000), "Invited Commentary: Particulate Matter-Mortality Exposure-Response Relations and Threshold," *American Journal of Epidemiology*, 152, 407–412.
- Pope, C. A., Dockery, D., and Schwartz, J. (1995a), "Review of Epidemiological Evidence of Health Effects of Particulate Air Pollution," *Inhalation Toxicology*, 47, 1–18.
- Pope, C. A., Thun, M., Namboodiri, M., Dockery, D., Evans, J., Speizer, F., and Heath, C. (1995b), "Particulate Air Pollution as a Predictor of Mortality in a Prospective Study of U.S. Adults," *American Journal of Respiratory and Critical Care Medicine*, 151, 669–674.
- Raftery, A., and Lewis, S. (1992), "How Many Iterations in the Gibbs Sampler?," in *Bayesian Statistics 4, Proceedings of the Fourth Valencia International Meeting*, pp. 763–774.
- Samet, J. M., Dominici, F., Curriero, F., Coursac, I., and Zeger, S. L. (2000a), "Fine Particulate Air Pollution and Mortality in 20 US Cities: 1987–1994" (with discussion), *New England Journal of Medicine*, 343, 1742–1757.
- Samet, J. M., Zeger, S. L., and Berhane, K. (1995), *The Association of Mortality and Particulate Air Pollution*, Cambridge, MA: Health Effects Institute.
- Samet, J. M., Zeger, S. L., Dominici, F., Curriero, F., Coursac, I., Dockery, D., Schwartz, J., and Zanobetti, A. (2000b), *The National Morbidity, Mortality, and Air Pollution Study (HEI Project No. 96-7): Morbidity and Mortality from Air Pollution in the United States*, Cambridge, MA: Health Effects Institute.
- Samet, J. M., Zeger, S. L., Dominici, F., Dockery, D., and Schwartz, J. (1999), *The National Morbidity, Mortality, and Air Pollution Study (HEI Project No. 96-7): Methods and Methodological Issues*, Cambridge, MA: Health Effects Institute.
- Samet, J. M., Zeger, S. L., Kelsall, J. E., Xu and Kalkstein, L. (1998), "Does Weather Confound or Modify the Association of Particulate Air Pollution with Mortality?" *Environmental Research*, 77, 9–19.
- Samet, J. M., Zeger, S. L., Kelsall, J., Xu, J., and Kalkstein, L. (1997), "Air Pollution, Weather and Mortality in Philadelphia," in *Particulate Air Pollution and Daily Mortality: Analyses of the Effects of Weather and Multiple Air Pollutants, The Phase IB report of the Particle Epidemiology Evaluation Project*, Cambridge, MA: Health Effects Institute.
- Schwartz, J. (1994), "Air Pollution and Daily Mortality: A Review and Meta Analysis," *Environmental Research*, 64, 36–52.
- Tanner, M. A. (1991), *Tools for Statistical Inference—Observed Data and Data Augmentation Methods*, Lecture Notes in Statistics, Vol. 67, New York: Springer-Verlag.
- Wakefield, J. C., Gelfand, A. E., and Smith, A. F. M. (1991), "Efficient Generation of Random Variates via the Ratio-of-Uniforms Method," *Statistics and Computing*, 1, 129–133.
- Zeger, S. L., Dominici, F., and Samet, J. M. (1999), "Harvesting-Resistant Estimates of Pollution Effects on Mortality," *Epidemiology*, 89, 171–175.



Annotations from j00125r1.pdf

Page 3

Annotation 1; Label: dominici; Date: 12/6/2001 2:45:41 AM
au: should these be 'df/yr' too?

Page 6

Annotation 1; Label: dominici; Date: 12/6/2001 2:46:43 AM
au: what is 'std' here? spell out on first use.

Page 9

Annotation 1; Label: dominici; Date: 12/6/2001 2:51:01 AM
au: spell out:?

Page 10

Annotation 1; Label: dominici; Date: 12/6/2001 2:52:31 AM
au: OK?

Page 12

Annotation 1; Label: dominici; Date: 12/6/2001 2:53:13 AM
au: can you update?

Annotation 2; Label: dominici; Date: 12/6/2001 2:53:38 AM
au: can you update?

Annotation 3; Label: dominici; Date: 12/6/2001 2:54:11 AM
au: pages?

Annotation 4; Label: dominici; Date: 12/6/2001 2:54:45 AM
au: Ed(s)? city and publisher?

Annotation 5; Label: dominici; Date: 12/6/2001 2:56:52 AM
au: I can't find text citations for American Thoracic Society and Bascomb(1996b) and Daniels and Kass(1998). Either add these or delete them from the ref. list.