

An informative Bayesian structural equation model to assess source-specific health effects of air pollution

MARGARET C. NIKOLOV*, BRENT A. COULL, PAUL J. CATALANO

*Department of Biostatistics, Harvard School of Public Health,
Boston, MA 02115, USA
meg.nikolov@gmail.com*

JOHN J. GODLESKI

*Department of Environmental Health, Harvard School of Public Health,
Boston, MA 02115, USA*

SUMMARY

A primary objective of current air pollution research is the assessment of health effects related to specific sources of air particles or particulate matter (PM). Quantifying source-specific risk is a challenge because most PM health studies do not directly observe the contributions of the pollution sources themselves. Instead, given knowledge of the chemical characteristics of known sources, investigators infer pollution source contributions via a source apportionment or multivariate receptor analysis applied to a large number of observed elemental concentrations. Although source apportionment methods are well established for exposure assessment, little work has been done to evaluate the appropriateness of characterizing unobservable sources thus in health effects analyses. In this article, we propose a structural equation framework to assess source-specific health effects using speciated elemental data. This approach corresponds to fitting a receptor model and the health outcome model jointly, such that inferences on the health effects account for the fact that uncertainty is associated with the source contributions. Since the structural equation model (SEM) typically involves a large number of parameters, for small-sample settings, we propose a fully Bayesian estimation approach that leverages historical exposure data from previous related exposure studies. We compare via simulation the performance of our approach in estimating source-specific health effects to that of 2 existing approaches, a tracer approach and a 2-stage approach. Simulation results suggest that the proposed informative Bayesian SEM is effective in eliminating the bias incurred by the 2 existing approaches, even when the number of exposures is limited. We employ the proposed methods in the analysis of a concentrator study investigating the association between ST-segment, a cardiovascular outcome, and major sources of Boston PM and discuss the implications of our findings with respect to the design of future PM concentrator studies.

Keywords: Factor analysis; Latent variables; Measurement error; Multivariate receptor model; Source apportionment.

*To whom correspondence should be addressed.

1. INTRODUCTION

Epidemiological studies have consistently demonstrated increased morbidity and mortality outcomes associated with elevated levels of airborne particulate matter (PM; Dockery *and others*, 1993; Dominici *and others*, 2002). The most direct way to reduce health effects of ambient air particles is to regulate the sources of pollution having adverse effects. In order to set protective standards, researchers must first establish the risk associated with the different pollution sources. Thus, one of the major objectives of current PM research is the assessment of the health effects related to specific sources of air pollution, such as power plants and motor vehicles. Recent research suggests that emissions from different sources exhibit differing levels of toxicity (Clarke *and others*, 2000; Laden *and others*, 2000; Godleski *and others*, 2002; Wellenius *and others*, 2003). However, quantifying source-specific risk is a challenge because current studies investigating the health effects of air pollution do not observe the PM contributions of the sources directly. Rather, exposures consist of samples of ambient air, or concentrated versions of ambient air, which reflect dynamic mixtures of source contributions. However, by taking into account the chemical fingerprints of known sources, an assessment of the chemical composition of exposure provides indirect information on the source contributions.

The exposure assessment literature contains an ample amount of research that focuses on estimation of source-specific contributions from a complex mixture of air pollution (for a review, see Hopke, 2003). Methods such as source apportionment and multivariate receptor modeling use factor analytic techniques to estimate the contributions of a small number of pollution sources from the measured mixture components, elements and other compounds. Although receptor modeling is well developed for exposure assessment, little work has been done to evaluate the appropriateness of characterizing unobservable sources in this way to estimate source-specific health effects. The problem can be thought of as an exposure measurement error problem (Carroll *and others*, 1995), whereby the PM exposure generated from a particular source is estimated rather than known or measured directly.

At present, existing source-specific health effects analyses rely on approaches that do not take into account the uncertainty associated with estimated source contributions. A “2-stage” strategy uses estimated source contributions from a factor analysis to assess the impacts of specific pollution sources on health effects (Laden *and others*, 2000; Clarke *and others*, 2000). A variation on this 2-stage strategy is the tracer approach, whereby the estimated source contributions are replaced by the elemental concentrations of a distinct set of tracers in the health effects analysis (Wellenius *and others*, 2003).

The statistical properties of the PM health effects estimates obtained from these existing approaches are not well understood. Previous statistical research has shown that, in simpler models, measurement error associated with estimated latent variables can lead to bias in the subsequent regression coefficient estimates (Tsiatis *and others*, 1995). It is unclear whether this bias will occur in PM research as the mixtures typically observed in PM exposures are quite different from those in other latent variable settings. As noted by Dominici *and others* (2003), there remains much work to be done in order to understand these estimates from a statistical standpoint and to assess the reliability of these estimates of association between pollution sources and health outcomes.

In this paper, we propose a structural equation framework for assessing source-specific health effects using speciated data in the form of elemental concentrations. This approach corresponds to jointly fitting a multivariate receptor model to the exposure data and a model for the health outcome given source contributions. Because the source contributions and health effects are modeled jointly, resulting inferences on the health effects account for the fact that uncertainty is associated with the exposures of interest.

This work is motivated by animal toxicology studies evaluating the mechanisms of morbidity and mortality associated with inhalation of concentrated air particles (CAPs) conducted at the Harvard School of Public Health (HSPH) (Godleski *and others*, 2000). Harvard researchers have implemented multiple animal toxicology studies to investigate the adverse effects of PM on cardiopulmonary and respiratory

activities in canines and rats. Samples of ambient Boston aerosol are collected and concentrated approximately 30 times by the Harvard ambient particle concentrator (Sioutas *and others*, 1995; Godleski *and others*, 2000) without altering the physical and chemical composition of the mixture. Animals are then exposed to the concentrated complex mixture for a given period of time, and cardiac and respiratory outcomes are monitored on each exposed animal. Because exposure is generated from ambient pollution, exposures are essentially random across days, and hence, a complete exposure assessment is made for each concentrated exposure. Data from these studies consist of the measured elemental concentrations of the concentrated air pollution mixture and the recorded health outcomes on the exposed animals. Because of the complexity of these studies, in any one study, investigators typically expose animals on approximately 20 unique exposure days.

The structural equation model (SEM), as well as the factor analysis model used in the 2-stage approach, typically involves a large number of parameters. Given the high dimensionality of the model, the typical study has an insufficient number of exposures to obtain reliable parameter estimates using maximum likelihood (ML). An approach for handling this problem is to consider a reduced number of elemental species for a health effects analysis (Clarke *and others*, 2000). To overcome the small-sample problem, we propose a fully Bayesian estimation approach that leverages historical exposure data from previous concentrator studies in defining informative priors on the parameters relating the measured exposures to the source contributions. This serves to pool exposure information from studies which are consistent in their collection and analysis of CAPs data.

The remainder of this paper is arranged as follows: Section 2 describes in detail the design and data from a study evaluating the effects of CAPs on myocardial ischemia in dogs (Wellenius *and others*, 2003). Section 3 presents the SEM, and Section 4 discusses the informative Bayesian approach to estimation. Section 5 presents a simulation study to examine the statistical properties of health effects estimates obtained with the tracer, 2-stage, and structural equation methodologies. Section 6 demonstrates an application of the informative Bayesian SEM to analyze the study of Wellenius *and others* (2003). Finally, in Section 7 we discuss our findings along with implications for the design of future PM concentrator studies.

2. DATA

Wellenius *and others* (2003) described results from a concentrator study examining the effects of inhaled CAPs on myocardial ischemia in dogs. The study subjects were 6 retired mongrel breeder dogs, each initially fitted with a balloon occluder around the left anterior descending coronary artery. The study design consisted of pairs of dogs undergoing 3 or 4 consecutive days of exposure and evaluation. On each day of the cycle, the dogs were put in side-by-side chambers and simultaneously underwent a continuous 6-h exposure to either CAPs or filtered air (Sham). The dogs were randomly assigned to CAPs exposure; one dog was exposed to CAPs on the second day of the cycle, while the other dog received CAPs exposure on the third day. Immediately following each exposure period, the dogs underwent a 5-min coronary artery occlusion and were monitored via continuous electrocardiogram (ECG). The primary outcome of interest was peak ST-segment elevation, a marker for myocardial ischemia. The study protocol was repeated multiple times for a total of 18 complete and successful exposure cycles. Table 1 summarizes the study design. This experiment was conducted according to the principles and regulations of the National Institutes of Health under protocols approved by the Harvard Medical Area Standing Committee on Animals.

Samples of the CAPs exposures were collected and analyzed. Each CAPs exposure was measured for sulfate (SULF) via ion chromatography, black carbon (BC) using an aethalometer, elemental carbon (EC) and organic carbon (OC) determined with a thermal and optical reflectance method, and elemental concentrations (in $\mu\text{g}/\text{m}^3$) collected via X-ray fluorescence, specifically: aluminum (Al), arsenic (As), barium (Ba), bromine (Br), calcium (Ca), cadmium (Cd), chlorine (Cl), chromium (Cr), copper (Cu),

Table 1. *Order and timing of exposures (Wellenius and others, 2003)*

Sequence	Dog	Start date	Day 1	Day 2	Day 3	Day 4
1	1	9/13/2000	Sham	Sham	CAPs	—
2	2	9/13/2000	Sham	CAPs	Sham	—
3	3	12/13/2000	Sham	CAPs	Sham	—
4	2	12/13/2000	Sham	Sham	CAPs	—
5	2	1/9/2001	Sham	Sham	CAPs	—
6	3	1/24/2001	Sham	Sham	CAPs	—
7	2	1/24/2001	Sham	CAPs	Sham	—
8	3	2/6/2001	Sham	CAPs	Sham	Sham
9	2	2/6/2001	Sham	Sham	CAPs	Sham
10	4	2/13/2001	Sham	CAPs	Sham	—
11	3	2/20/2001	Sham	Sham	CAPs	Sham
12	2	2/20/2001	Sham	CAPs	Sham	Sham
13	5	2/27/2001	Sham	CAPs	Sham	Sham
14	6	2/27/2001	Sham	Sham	CAPs	Sham
15	3	3/7/2001	Sham	CAPs	Sham	—
16	5	3/12/2001	Sham	Sham	CAPs	Sham
17	6	3/12/2001	Sham	CAPs	Sham	Sham
18	6	3/27/2001	Sham	CAPs	Sham	Sham

Note that 4 pairings were not complete due to failed sequences.

iron (Fe), potassium (K), manganese (Mn), nickel (Ni), sodium (Na), lead (Pb), sulfur (S), selenium (Se), silicon (Si), titanium (Ti), vanadium (V), and zinc (Zn). The Sham exposures were assumed to have zero concentration of all elements, as this had been confirmed in earlier test runs of the concentrator (Sioutas and others, 1995; Lawrence and others, 2004).

Extensive research into the composition of Boston aerosol (Oh and others, 1997; Clarke and others, 2000; Batalha and others, 2002) has revealed 4 major sources of PM pollution in this area; resuspended road dust consisting mainly of the crustal elements (Si and Al), coal-fired power plants (S and SULF), oil combustion primarily for home heating (Ni and V), and motor vehicle exhaust (BC, OC, and EC). Wellenius and others (2003) considered linear mixed models for log transformed peak ST-segment, using individual elemental concentrations as tracer representatives of known sources of air pollution in Boston. These authors chose silicon, sulfur, nickel, and BC to represent road dust, power plants, oil combustion, and motor vehicles, respectively, and found a strong positive association between log peak ST-segment and resuspended road dust. A question that naturally arises is whether measurement error associated with tracer representatives of the source contributions obscured the relationship between log peak ST-segment and the other pollution sources. In this article, we analyze the data using methods that account for the uncertainty in estimated source contributions.

3. MODEL AND NOTATION

3.1 Modeling framework

We propose a full-likelihood approach that estimates the health effects by fitting the receptor and health outcome models jointly. The joint model, which falls within the SEM framework (Bollen, 1989), is

$$\mathbf{X}_t = \mathbf{\Lambda}\boldsymbol{\eta}_t + \boldsymbol{\epsilon}_t^X, \quad (3.1)$$

$$Y_t = \alpha + \boldsymbol{\beta}^T \boldsymbol{\eta}_t + \epsilon_t^Y, \quad (3.2)$$

where for a given time t , \mathbf{X}_t is the vector of p elemental concentrations, $\boldsymbol{\eta}_t$ is the vector of the k unobserved source contributions, and Y_t is the health outcome. We assume that Y_t represents a single continuous variable and that k is known. The model for \mathbf{X}_t is the factor analysis model for the exposure analysis (Park, 2001), where $\mathbf{\Lambda}$ is the $(p \times k)$ matrix of factor loadings, also known as the factor pattern, and $\boldsymbol{\epsilon}_t^X \stackrel{\text{iid}}{\sim} N_p(\mathbf{0}, \boldsymbol{\Psi})$ for diagonal $\boldsymbol{\Psi}$. The vector $(\lambda_{1j}, \lambda_{2j}, \dots, \lambda_{pj})$ may be viewed as the profile of pollution source j . The parameters $\boldsymbol{\beta}$ quantify the k source-specific health effects, and $\boldsymbol{\epsilon}_t^Y \stackrel{\text{iid}}{\sim} N(0, \sigma_Y^2)$. Standard factor analysis assumes $\boldsymbol{\eta}_t \stackrel{\text{iid}}{\sim} N_k(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

3.2 Distributional assumptions

We extend the standard SEM defined by (3.1)–(3.2) in 2 ways. First, we truncate a normal distribution for $\boldsymbol{\eta}$ to ensure the physical nonnegativity of source contributions

$$\boldsymbol{\eta} \stackrel{\text{iid}}{\sim} N_k(\boldsymbol{\mu}, \boldsymbol{\Sigma}) I(\boldsymbol{\eta} \geq \mathbf{0}).$$

Alternatively, we could specify a lognormal distribution. In our application, we assess the sensitivity of our conclusions to distributional assumptions on $\boldsymbol{\eta}$ by fitting the SEM both ways. The specification of nonnegative source contributions extends the standard factor analysis, which does not restrict the domain of $\boldsymbol{\eta}$ in the model and allows negative source contributions. Positive matrix factorization (Paatero and Tapper, 1994) is an alternative method that uses constrained weighted least squares to ensure nonnegative source contributions.

Second, to accommodate the repeated measures design of these studies, we build random effects into the model for the health outcome

$$Y_{st} = \alpha + \boldsymbol{\beta}^T \boldsymbol{\eta}_t + \mathbf{Z}_{st}^T \mathbf{b}_s + \boldsymbol{\epsilon}_{st}^Y,$$

where Y_{st} is the health outcome, \mathbf{Z}_{st} is the vector of covariates for unit s at time t , \mathbf{b}_s is the vector of random effects for unit s , $\mathbf{b}_s \stackrel{\text{iid}}{\sim} N(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{b}})$, $\boldsymbol{\epsilon}_{st}^Y \stackrel{\text{iid}}{\sim} N(0, \sigma_Y^2)$, and $\mathbf{b} \perp \boldsymbol{\epsilon}^Y$.

3.3 Model identifiability

The SEM defined by (3.1)–(3.2) is not identifiable without further assumptions. Because the source profiles are unknown and the source contributions are unobserved, the SEM does not have a unique solution. However, the model may be made identifiable by constraining parameters in $\mathbf{\Lambda}$. We consider the following 2 sets of identifiability conditions, which result in a confirmatory factor analysis (Park and others, 2002):

- C1: There are at least $k - 1$ zero elements in each column of $\mathbf{\Lambda}$.
- C2: The rank of $\mathbf{\Lambda}^{(j)}$ is $k - 1$, where $\mathbf{\Lambda}^{(j)}$ is the matrix composed of the rows containing the assigned 0's in the j th column with those assigned 0's deleted.
- C3: $\lambda_{ij} = 1$ for some i ($i = 1, 2, \dots, p$) for each $j = 1, 2, \dots, k$.
- D1: There are at least k rows in $\mathbf{\Lambda}$ with each of the k rows containing only one nonzero element.
- D2: Same as C2.
- D3: Same as C3.

The C1–C3 conditions assume that there are at least $k - 1$ elements per source that are not associated with that source and that this group of elements is not the same for all sources. The D1–D3 conditions

assume that there is at least one “tracer” element for each source in the sense that the tracer does not load on other sources. Further, as noted by Park *and others* (2001), these conditions identify the loadings up to normalization. Thus, for each source, we specify one loading to be equal to 1, effectively placing the source contribution on the scale of the element having the constrained loading of 1 for that source. Combining all these constraints, the D1–D3 conditions effectively use the concept of tracers in the factor analysis model but for each source supplement this tracer information with that contained in nontracer elements.

The C1–C3 conditions and the D1–D3 conditions are each sufficient but not necessary to establish identifiability. While there exist alternative conditions, other commonly used proposals are also sufficient but not necessary. For instance, Park *and others* (2002) proposed sufficient conditions which, instead of placing constraints on the factor loadings, assume that some sources are absent on some days. These authors argued that in some settings, this alternative set of constraints may be plausible if one knows that a particular source, such as a power plant in the region, has been shut down for some period of time. In the same vein, Bandeen-Roche (1994) considered situations in which a subset of the source contributions is known. In our setting, however, we do not have information on the presence or absence of a particular source on a particular day. Thus, given the existing literature on the pollution mixture in the Boston area (Oh *and others*, 1997), it seems safer to assume that certain elements are not markers for certain sources. Bollen (1989) also gave some rules that help to determine whether a model is identifiable, but, as noted by this author, these rules are also either necessary or sufficient, but not both.

4. ESTIMATION

Standard SEMs may be fit via ML using existing latent variable software, such as Mplus (Muthen and Muthen, 1998) or the “sem” package in R (R Development Core Team, 2003). Due to the large number of parameters involved, these methods require a large sample size for asymptotic optimality of the resulting ML estimators. In studies motivating this research, the number of exposures rarely exceeds 20 and ML methods break down. The problem is due to the large number of parameters in the factor analysis model (3.1). Twenty days of exposure are simply too few to obtain a reliable receptor model fit in most cases.

To overcome the small number of unique exposure days, we propose an informative Bayesian approach to model fitting. This approach is especially appealing considering that HSPH researchers have conducted multiple concentrator studies, all of which are consistent in their collection and analysis of exposure data. It is reasonable to pool the exposure data from prior studies to estimate the profiles of PM sources in Boston. An informative Bayesian approach leverages historical exposure data to obtain more reliable estimates of the source profiles, thus improving our ability to estimate the health effects investigated in an individual study.

The Bayesian approach incorporates information from previous studies through specification of the priors. In this case, a preliminary factor analysis of the historical exposure data provides prior information on the unknown factor pattern $\mathbf{\Lambda}$. Let $\boldsymbol{\lambda}$ be the $k \times (p - k)$ vector of all unconstrained factor loadings, $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_1^T, \boldsymbol{\lambda}_2^T, \dots, \boldsymbol{\lambda}_k^T)^T$, where $\boldsymbol{\lambda}_j$ is the vector of unconstrained loadings for source j , $j = 1, \dots, k$. Let $\widehat{\boldsymbol{\lambda}}^{(\text{hist})}$ and $\widehat{\text{Var}}(\widehat{\boldsymbol{\lambda}}^{(\text{hist})})$ represent the posterior mean and covariance of the factor loadings obtained from a Bayesian factor analysis of the historical data. For the informative Bayesian SEM, the prior distribution on the free parameters in $\mathbf{\Lambda}$ may be defined as

$$\boldsymbol{\lambda} \sim N_{k(p-k)}(\widehat{\boldsymbol{\lambda}}^{(\text{hist})}, \widehat{\text{Var}}(\widehat{\boldsymbol{\lambda}}^{(\text{hist})})), \quad (4.1)$$

while the constrained loadings are treated as fixed constants in the likelihood.

5. SIMULATION STUDY

We conducted a simulation study to examine the statistical properties of the health effects estimates obtained via the tracer, 2-stage, and structural equation approaches. In the interest of direct comparison between the various approaches, we assume a normal mean zero distribution on the source contributions to ensure that our assessments are not confounded by distributional assumptions made by different implementations of SEMs.

In order to make our findings most relevant to the HSPH concentrator studies, we based our simulations on the known sources of Boston PM pollution described previously. We obtained realistic parameter settings for Λ , Σ , and Ψ from a confirmatory factor analysis of the complete aggregated exposure data ($n = 178$). We conducted our analysis on a subset of $p = 13$ elements deemed to be major components of the 4 known sources of Boston PM; silicon (Si), sulfur (S), nickel (Ni), OC, aluminum (Al), titanium (Ti), calcium (Ca), SULF, selenium (Se), vanadium (V), bromine (Br), BC, and EC. Since convergence problems are common when elemental concentrations are on widely different scales, each element was scaled by its sample standard deviation, which is equivalent to conducting a factor analysis on the sample correlation matrix, as opposed to the sample covariance matrix.

We constrained one “tracer” element for each of the $k = 4$ sources according to the D1–D3 identifiability conditions. We chose silicon, sulfur, nickel, and OC to identify road dust, power plants, oil combustion, and motor vehicles, respectively. A preliminary exploratory factor analysis justified the “tracer” identifiability conditions, since the estimated factor loadings of silicon, sulfur, nickel, and OC were low (<0.2) on all but a single source. Due to space constraints, values of the parameter settings for Λ , Σ , and Ψ may be found in an accompanying online technical report (Nikolov *and others*, 2006). Settings for the health model parameters were motivated by the investigation of Wellenius *and others* (2003) of PM effects on heart rate; $\alpha = 86$, $\beta = 2$, and $\sigma_Y = 8$, yielding an effect size of $\delta = 0.25$.

To simulate exposure, we generated source contributions from $\eta \sim N_4(\mathbf{0}, \Sigma)$ and elemental concentrations from $\mathbf{X}|\eta \sim N_{13}(\Lambda\eta, \Psi)$. We generated health outcomes assuming a health effect from a single source. Specifically, for a given source j , the health outcome was simulated from a simple linear regression model assuming only an effect of that source. For example, \mathbf{y}_1 is a vector of simulated health outcomes where the health effect is associated with the first factor, road dust. For each set of generated exposures, we generate 4 sets of health outcomes, \mathbf{y}_1 , \mathbf{y}_2 , \mathbf{y}_3 , and \mathbf{y}_4 , where the health effect corresponds to the different pollution sources, road dust, power plants, oil combustion, and motor vehicles, respectively. We analyzed these simulated health outcomes separately.

Exposures, source contributions and elemental concentrations, were generated for $n \in \{20, 100\}$ days. Although the HSPH concentrated particle experiments typically do not run with 100 exposure days, we included this hypothetical scenario to confirm that any deficiencies of the ML SEM are due to a small number of exposure days. The health outcomes were generated for 2 animals per exposure day, for a total of $2n \in \{40, 200\}$ outcomes.

We obtained health effects estimates using 5 different strategies:

1. Known source contributions: Although source contributions are not directly measured in the studies motivating this research, here we simulate them so that they are effectively known. We estimate the health effects using these source contributions as predictors in the health effects model. This setting represents an unobtainable “gold standard” and is provided for reference only.
2. Tracer approach: We estimated the health effects using the elemental concentrations of the distinct set of the tracers, silicon, sulfur, nickel, and OC, as covariates in the health effects model.
3. 2-stage approach: We first conducted a confirmatory factor analysis on all simulated elements, constraining the factor loadings for the tracer elements, silicon, sulfur, nickel, and OC, according to the D1–D3 identifiability conditions. We then fit the health effects model using the estimated source contributions as predictors in the health model.

4. ML SEM: We estimated the receptor and health effects models (3.1) and (3.2) jointly using ML in Mplus (Muthen and Muthen, 1998). This approach imposed the D1–D3 identifiability conditions but did not use any historical exposure information.
5. Bayesian SEM: We estimated the receptor and health effects models jointly using an informative Bayesian approach. To obtain informative priors on the source profiles, we conducted a confirmatory factor analysis on a single simulated historical data set of $n = 200$ exposures, based on the D1–D3 identifiability constraints and tracers, silicon, sulfur, nickel, and OC. We defined informative priors on the source profiles using (4.1) and set vague priors on the remaining parameters; $IG(0.01, 0.01)$ on $\{\Sigma_{jj}\}$, $\{\Psi_{ii}\}$, and σ_Y^2 and $N(0, 1000)$ on α and $\{\beta_j\}$. The Bayesian SEM was fit using the Markov chain Monte Carlo (MCMC) method in WinBUGS (Spiegelhalter *and others*, 2000); for each fit, we ran 25 000 iterations, discarding 20 000 as burn-in and thinning by 5, for a total of 1000 posterior samples for estimation and inference. We randomly checked convergence on multiple simulated data sets and saw evidence of good mixing and convergence in every case.

We ran 500 simulations to assess the statistical properties of the health effects estimates of a typical concentrator study with $n = 20$ days of exposure and an additional 500 simulations to evaluate the estimates of a hypothetical study with $n = 100$ exposure days.

Table 2 displays the health effects estimates and corresponding simulation standard errors, obtained with the 5 different methodologies. Although the health effects were estimated with a model that included terms to represent all 4 sources, the table presents only the estimate for the source on which the health effect was simulated. For example, the first column contains the health effects estimates corresponding to road dust, since this column reflects the analysis of the y_1 outcome, where the health effect was simulated on the road dust source. The estimated coefficients for the other 3 sources, power plants, oil combustion, and motor vehicles, were always all approximately zero and, hence, are not included in the tables. In all cases, our estimates of the null coefficients were unbiased, and therefore, we display only the estimates for which the truth is $\beta = 2$.

In Table 2, we see that the health effects estimates obtained with the tracer approach demonstrate the typical attenuation of effect associated with measurement error in this simple setting. This finding is consistent in the simulation study based on $n = 100$ exposure days as well (results not shown due to space constraints). In fact, we can calculate the attenuation factor γ associated with each tracer estimate, since we know the amount of measurement error associated with each of the tracer elements, quantified by ψ_{Si} , ψ_S , ψ_{Ni} , and ψ_{OC} . Because our simulations are based on a factor pattern with a unique tracer for each source, uncorrelated factors, and normality, for a given source j , $\gamma_j = (\Sigma_{jj})/(\Sigma_{jj} + \Psi_{jj})$. Here, the attenuation factors are 0.97, 0.97, 0.87, and 0.78 for the road dust, power plants, oil combustion, and motor vehicles effects, respectively. In our simulation study, we are able to correct for the bias induced by measurement error and obtain reliable health effects estimates using the tracer approach. However, in practical settings, the variance parameters $\{\Sigma_{jj}\}$ and $\{\Psi_{jj}\}$ are typically unknown, and the small number of exposure days prevents us from obtaining reliable estimates of these correction factors.

Table 2. *Simulation study: health effect estimates (standard errors) for $n = 20$*

Method	Road dust	Power plants	Oil combustion	Motor vehicles
Known η	1.94 (0.04)	2.04 (0.05)	1.99 (0.05)	1.99 (0.05)
Tracer	1.85 (0.04)	1.99 (0.05)	1.74 (0.05)	1.58 (0.05)
2-stage	1.91 (0.04)	1.93 (0.05)	1.66 (0.05)	1.83 (0.06)
ML SEM	1.92 (0.04)	1.97 (0.05)	1.77 (0.06)	1.89 (0.06)
Bayes SEM	1.90 (0.04)	2.02 (0.05)	1.91 (0.06)	1.95 (0.05)

Alternatively, the 2-stage approach amounts to estimating the correction terms and adjusting the health effects estimates accordingly by using $\hat{\eta} = \hat{E}(\eta|\mathbf{Y})$ in the health outcome model. In this way, the 2-stage approach may be viewed as a form of regression calibration (Carroll *and others*, 1995). The simulation study demonstrates attenuation in the 2-stage health effects estimates based on $n = 20$ exposure days; however, we attribute this bias to the small number of exposures. In the study based on $n = 20$ exposures, the receptor model failed to converge in 19 out of 500 (3.8%) simulations. However, in the study based on $n = 100$ days of exposure, all receptor models converged. Furthermore, the 2-stage estimates based on $n = 100$ exposure days are all very similar to the estimates obtained with the known source contributions, and all estimates are within twice the simulation standard error of the truth.

The SEM approach appears to offer a clear advantage to the tracer and 2-stage approaches, particularly in the case of a small number of exposures. However, in this small-sample context, the SEM estimates are distinguished by the method of estimation. The health effects estimates obtained with the informative Bayesian SEM are most similar to those obtained with known source contributions and are within twice the simulation standard error of the truth in almost all cases for $n = 20$. In contrast, the estimates obtained via ML appear to be biased downward, and the ML SEM estimate for oil combustion is well beyond twice the simulation standard error from the truth. As in the case of the 2-stage approach, we attribute these deficiencies in the ML SEM to the small number of exposures. In the study based on $n = 20$ exposure days, the ML method failed to converge in approximately 4% of the simulations. (The numbers of failures are 22, 22, 20, and 21 for the analysis of \mathbf{y}_1 , \mathbf{y}_2 , \mathbf{y}_3 , and \mathbf{y}_4 , respectively.) However, when we increase the number of exposures to $n = 100$ (results not shown), all 500 simulations converged, and the ML SEM performs almost exactly the same as the gold standard that uses known η .

Table 3 provides the estimated power, defined as the proportion of 95% confidence (or credible, for the Bayesian approaches) intervals that do not contain the null value of zero when $\beta_{\text{HE}} = 2$, as well as the estimated size, the proportion of 95% confidence (credible) intervals that do not contain a true value of $\beta_{\text{HE}}^* = 0$, for the study based on $n = 20$ exposure days. In this small study context, the 2-stage approach and the ML SEM exceed the expected size of 0.05 in all cases, indicating that these approaches are too liberal when the number of exposures is limited. For methods of approximately the same size (excluding the 2-stage and ML SEM), the Bayesian SEM is comparable to the “gold standard” based on known source contributions and has virtually the same sensitivity for detecting a true effect as the tracer approach. Likewise, in the study based on $n = 100$ exposure days (results not shown), the 2-stage approach and ML SEM are of the appropriate size (≈ 0.05), and all methods are very powerful ($>95\%$) at detecting a true effect in this setting.

Finally, we conducted an additional simulation study designed to investigate the impact of choosing incorrect identifiability constraints. One could argue that setting a single factor loading to zero when it is really greater than zero should have little impact on the resulting health effect estimates, whereas incorrectly setting $k \times (k - 1)$ loadings equal to zero may collectively have a larger effect. To check the impact of this misspecification, we conducted a simulation study assuming all loadings were nonzero, replacing the zero loadings in the previous simulation study to randomly generated values from a

Table 3. *Simulation study: power (size) analysis for $n = 20$*

Method	Road dust	Power plants	Oil combustion	Motor vehicles
Known η	54% (0.056)	44% (0.047)	40% (0.049)	40% (0.039)
Tracer	49% (0.050)	43% (0.045)	34% (0.053)	35% (0.039)
2-stage	53% (0.095)	42% (0.081)	36% (0.097)	34% (0.087)
ML SEM	62% (0.133)	53% (0.129)	47% (0.133)	46% (0.122)
Bayes SEM	50% (0.052)	44% (0.051)	35% (0.050)	37% (0.040)

Table 4. *Robustness study: health effect estimates (standard errors) for $n = 20$*

Method	Road dust	Power plants	Oil combustion	Motor vehicles
Known η	2.02 (0.07)	1.94 (0.05)	2.04 (0.05)	2.10 (0.05)
Tracer	1.96 (0.07)	1.90 (0.04)	1.81 (0.05)	1.65 (0.05)
Bayes SEM (D1–D3)	2.09 (0.07)	1.98 (0.05)	1.96 (0.06)	2.03 (0.05)
Bayes SEM (C1–C3)	2.11 (0.07)	1.96 (0.05)	2.00 (0.06)	2.05 (0.05)

Table 5. *Robustness study: power (size) analysis for $n = 20$*

Method	Road dust	Power plants	Oil combustion	Motor vehicles
Known η	54% (0.047)	41% (0.052)	40% (0.051)	45% (0.059)
Tracer	51% (0.053)	38% (0.061)	34% (0.054)	35% (0.062)
Bayes SEM (D1–D3)	49% (0.045)	41% (0.066)	36% (0.057)	42% (0.058)
Bayes SEM (C1–C3)	50% (0.047)	40% (0.063)	38% (0.057)	42% (0.060)

uniform(0, 0.2) distribution. For each simulated data set, we estimated the health effects using known source contributions, the tracer approach, the informative Bayesian SEM fit with the D1–D3 conditions, and the informative Bayesian SEM fit with the less restrictive C1–C3 identifiability conditions. Table 4 presents the source-specific health effects estimates and standard errors. Table 5 provides the corresponding power and size. These results are consistent with the results from the previous simulations for $n = 20$ exposure days, in that the Bayesian SEM approaches yield estimates similar to those obtained if the true source contributions are known and tests of the appropriate size. Thus, the second study suggests that the proposed identifiability constraints do not have a large impact on inference as long as the loadings set to zero are not much larger than 0.2.

6. DATA ANALYSIS

6.1 Joint model

In this section, we implement our informative Bayesian SEM to analyze the source-specific PM health effects on myocardial ischemia in dogs (Wellenius *and others*, 2003). Analyses of the Wellenius data did not detect any pairing or period (day) effects but did suggest there may be a carryover effect from the CAPs exposure. As a result, the authors excluded Sham exposures following the CAPs exposure in their analyses. The final analysis was based on a total of 43 measured health outcomes, corresponding to 18 CAPs and 25 Sham exposures.

Wellenius *and others* (2003) found a large amount of variability within each dog by cycle combination and therefore included a random effect for sequence (see Table 1). The authors estimated the source-specific health effects of PM with the following model:

$$Y_{td} = \alpha + \boldsymbol{\beta}^T \mathbf{x}_{td}^{(T)} + b_t + \epsilon_{td}^Y,$$

where Y_{td} is log(peak ST-segment), $\mathbf{x}_{td}^{(T)}$ is the vector of elemental concentrations for silicon, sulfur, nickel, and BC for sequence t and day d , b_t is the random sequence effect, $b_t \stackrel{\text{iid}}{\sim} N(0, \sigma_b^2)$, $\epsilon_{td}^Y \stackrel{\text{iid}}{\sim} N(0, \sigma_Y^2)$, and $b \perp \epsilon^Y$.

Accordingly, we fit the following informative SEM model:

$$\mathbf{X}_t = \mathbf{\Lambda}\boldsymbol{\eta}_t + \boldsymbol{\epsilon}_t^{\mathbf{X}},$$

$$Y_{td} = \alpha + \boldsymbol{\beta}^T (I_{\text{CAPS}_{td}} \times \boldsymbol{\eta}_t) + b_t + \epsilon_{td}^Y,$$

where $\boldsymbol{\epsilon}_t^{\mathbf{X}} \stackrel{\text{iid}}{\sim} N_p(\mathbf{0}, \boldsymbol{\Psi})$ for diagonal $\boldsymbol{\Psi}$, $b_t \stackrel{\text{iid}}{\sim} N(0, \sigma_b^2)$, $\epsilon_{td}^Y \stackrel{\text{iid}}{\sim} N(0, \sigma_Y^2)$, and $b \perp \epsilon^Y$. $I_{\text{CAPS}_{td}}$ is an indicator of CAPs exposure on day d in sequence t . This indicator provides for the Sham exposures and operates on the assumption that all source contributions are null in filtered air, that is, if the dog in sequence t is exposed to Sham on day d , $I_{\text{CAPS}_{td}} = 0$ and $Y_{td} = \alpha + b_t + \epsilon_{td}^Y$. Based on this specification, the receptor model applies to the CAPs exposures only, while the health effects regression is fit on all outcomes. Finally, to respect the nonnegativity of source contributions, we truncate the normal distribution on the latent variables, $\boldsymbol{\eta}_t \stackrel{\text{iid}}{\sim} N_k(\boldsymbol{\mu}, \boldsymbol{\Sigma})I(\boldsymbol{\eta}_t \geq \mathbf{0})$. As noted in Section 3.2, we also fit the model assuming log-normal source contributions to assess the sensitivity of our conclusions to distributional assumption on $\boldsymbol{\eta}$.

6.2 Prior elicitation

To obtain prior information on the source profiles, we fit a Bayesian confirmatory factor analysis on the scaled historical data ($n = 160$). We assumed the $k = 4$ major sources of Boston PM, and we assumed that these sources are independent. To be consistent with the Wellenius *and others* (2003) analysis, here we chose silicon to identify road dust, sulfur for power plants, nickel for oil combustion, and BC for motor vehicles.

It is thought that BC is a better marker of motor vehicles than OC. However, exploratory factor analyses consistently estimated moderate loadings for BC on several factors. Therefore, in our analysis, we apply the more flexible C1–C3 identifiability conditions described in Section 3.2. According to these conditions, we need to constrain $k - 1 = 3$ loadings to zero on each profile, while ensuring a distinct set of constraints for each source. In order to set meaningful constraints, we consulted our exploratory results and identified distinct sets of three near zero (< 0.2) loadings per factor. We constrain to zero the following factor loadings: sulfur, nickel, and EC on road dust; silicon, vanadium, and OC on power plants; aluminum, SULF, and OC on oil combustion; and aluminum, SULF, and nickel on motor vehicles.

We fit the confirmatory factor analysis to the historical data using MCMC in WinBUGS. We set vague priors on all parameters, specifying $\text{IG}(0.01, 0.01)$ on $\{\boldsymbol{\Sigma}_{jj}\}$ and $\{\boldsymbol{\Psi}_{ii}\}$, $N(0, 1000)$ on $\{\mu_j\}$, and, following Park *and others* (2001), we truncate normal distributions for the unconstrained $\{\lambda_{ij}\}$, $\lambda_{ij} \sim N(0, 10000)I(\lambda_{ij} \geq 0)$, since negative components of source profiles are not interpretable. We ran 25 000 iterations, discarding 20 000 as burn-in and thinning by 5, for a total of 1000 posterior samples for estimation. Evaluation of autocorrelation and trace plots supported convergence.

Table 6 provides the posterior source profiles along with the estimated standard errors. Figure 1 displays the posterior correlations between the unconstrained factor loadings in 2 ways. The frequency distribution demonstrates that most of the posterior correlations are low ($-0.2 < \hat{\rho} < 0.2$), with the absolute values of a few correlations as large as 0.6. The temperature grid arranges these correlations by chemical species. This plot reveals that the few moderate correlations correspond primarily to pairs of loadings on the same element. Overall, the standard errors and correlations of the unconstrained factor loadings suggest that the additional exposure information provided by the nontracer elements makes a substantial contribution to the differentiation of the pollution sources.

6.3 Health effects analysis

We fit the informative Bayesian SEM to the data of Wellenius *and others* (2003) using MCMC in WinBUGS. We defined informative priors on the source profiles based on the posterior means in Table 6

Table 6. *Posterior source profiles: estimated factor loadings (SEs)*

Element	$\hat{\lambda}_{RD}$ (SE)	$\hat{\lambda}_{PP}$ (SE)	$\hat{\lambda}_{OC}$ (SE)	$\hat{\lambda}_{MV}$ (SE)
Si	1 [†] (—)	0 [†] (—)	0.05 (0.0019)	0.26 (0.0035)
S	0 [†] (—)	1 [†] (—)	0.02 (0.0009)	0.04 (0.0019)
Ni	0 [†] (—)	0.03 (0.0016)	1 [†] (—)	0 [†] (—)
OC	0.04 (0.0024)	0 [†] (—)	0 [†] (—)	1.48 (0.0098)
Al	1.00 (0.0018)	0.01 (0.0006)	0 [†] (—)	0 [†] (—)
Ti	0.96 (0.0029)	0.02 (0.0017)	0.18 (0.0036)	0.51 (0.0065)
Ca	0.95 (0.0035)	0.02 (0.0011)	0.21 (0.0042)	0.32 (0.0071)
SULF	0.01 (0.0004)	0.99 (0.0013)	0 [†] (—)	0 [†] (—)
Se	0.02 (0.0011)	0.66 (0.0048)	0.04 (0.0026)	0.38 (0.0079)
V	0.09 (0.0026)	0 [†] (—)	1.04 (0.0044)	0.04 (0.0031)
Br	0.10 (0.0032)	0.54 (0.0051)	0.18 (0.0040)	0.81 (0.0089)
BC	0.04 (0.0019)	0.44 (0.0039)	0.42 (0.0029)	1 [†] (—)
EC	0 [†] (—)	0.29 (0.0045)	0.42 (0.0034)	1.28 (0.0084)

[†]Denotes constrained parameters. RD, road dust; PP, power plants; OC, oil combustion; MV, motor vehicles; SE, standard error.

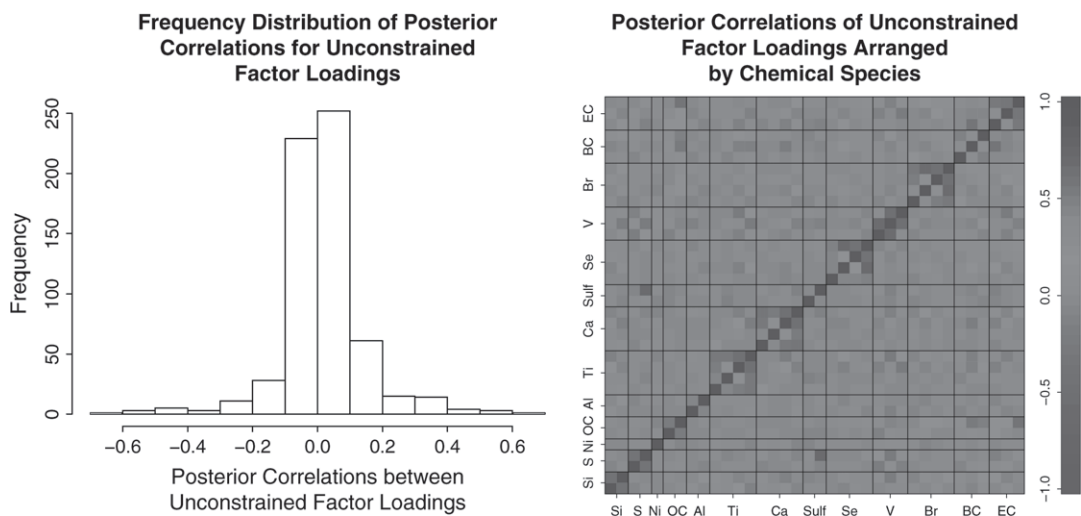


Fig. 1. Posterior correlations in λ from a Bayesian factor analysis of the historical exposure data.

and posterior covariances from the Bayesian factor analysis of the historical data. Here again, we truncated the normal distribution for the factor loadings, $\lambda \sim N_{k(p-k)}(\hat{\lambda}^{(hist)}, \widehat{\text{Var}}(\hat{\lambda}^{(hist)}))I(\lambda \geq \mathbf{0})$. We set vague priors on all other parameters, $\text{IG}(0.01, 0.01)$ on $\{\Sigma_{jj}\}$, $\{\Psi_{ii}\}$, σ_Y^2 , and σ_b^2 and $N(0, 1000)$ on $\{\mu_j\}$, α , and $\{\beta_j\}$. We ran 25 000 iterations, discarding 20 000 as burn-in and thinning by 5, for a total of 1000 posterior samples. We examined diagnostic trace and autocorrelation plots and found satisfactory convergence.

Table 7 displays the posterior means, standard errors, and 95% credible (confidence) intervals for the source-specific health effects estimated with the informative Bayesian SEM and the tracer approach.

Table 7. Comparative data analysis results

Method	Pollution source	Source tracer	$\hat{\beta}$	SE($\hat{\beta}$)	95% CI
Informative Bayesian SEM	Road dust	Si	0.154 [†]	0.063	(0.030, 0.276)
	Power plants	S	-0.071	0.072	(-0.217, 0.072)
	Oil combustion	Ni	-0.034	0.071	(-0.167, 0.115)
	Motor vehicles	BC	0.062	0.141	(-0.217, 0.341)
Source tracer approach	Road dust	Si	0.137 [†]	0.048	(0.037, 0.237)
	Power plants	S	-0.063	0.079	(-0.227, 0.100)
	Oil combustion	Ni	-0.036	0.069	(-0.178, 0.107)
	Motor vehicles	BC	0.014	0.088	(-0.170, 0.198)

[†]Denotes significance at $\alpha = 0.05$. SE, standard error.

In the SEM analysis, for each pollution source, the health effect estimate is on the scale of the element whose factor loading is constrained to 1. For example, we interpret the health effect estimate of road dust as the change in log peak ST-segment associated with an increase in the “contribution of road dust” on the scale of one standard deviation increase in the concentration of silicon. However, in the tracer analysis, we interpret the health effects estimates as the change in log peak ST-segment associated with one standard deviation increase in the concentration of corresponding tracer. For instance, here we interpret the health effect estimate of road dust as the change in log peak ST-segment associated with one standard deviation increase in the “concentration of silicon.” The tracer estimates provided in Table 7 do not precisely correspond to those reported in Wellenius *and others* (2003), which were based on ECG recordings from 2 precordial leads. Since readings from the 2 leads were highly correlated ($r > 0.8$), we restrict our analysis to log peak ST-segment recorded on a single lead (V5).

In accordance with Wellenius *and others* (2003), we found a significant effect of resuspended road dust on myocardial ischemia in dogs. Implementing our informative Bayesian SEM, we estimated an increase of 0.154 in log peak ST-segment associated with an increase in road dust contribution on the scale of one standard deviation of silicon, with 95% credible interval (0.030, 0.276). We did not find a significant change in log peak-ST segment associated with pollution from coal-fired power plants, oil combustion for home heating, or motor vehicle exhaust.

Thus, the informative Bayesian structural equation results suggest that the conclusions in the original analysis were not driven by unequal amounts of measurement error associated with the tracer representations of the 4 pollution sources. Given that the SEM analysis yields an estimate of the road dust coefficient similar to the silicon coefficient in the tracer analysis, one might wonder whether the more complex analysis was worth the effort. Then again, the SEM analysis does more than reinforce the significance of the road dust effect; it also reaffirms the lack of evidence of an association between this cardiac outcome and pollution from the other sources, in particular motor vehicles. Not surprisingly, given the results of our simulation study, the estimated health effect estimate for motor vehicles, $\hat{\beta}_{MV} = 0.062$, is more than 4 times the corresponding attenuated tracer estimate for BC, $\hat{\beta}_{BC} = 0.014$; however, this estimated association remains insignificant. Confirmation of a nonsignificant effect of motor vehicles is equally important in this setting.

To ensure that our conclusions are not sensitive to distributional assumptions, we reran the informative Bayesian SEM under the specification of lognormal source contributions. We estimated $\hat{\beta}_{RD} = 0.164$, $\hat{\beta}_{PP} = -0.066$, $\hat{\beta}_{OC} = -0.030$, and $\hat{\beta}_{MV} = 0.026$ and detected a significant effect of road dust only. Thus, both sets of distributional assumptions yield findings that agree with those in Wellenius *and others* (2003).

7. DISCUSSION

In this paper, we considered methods to assess source-specific health effects of complex mixtures of PM pollution. One objective was to evaluate the statistical properties of estimates obtained with methods currently used in practice, the tracer approach and the 2-stage approach, for multivariate pollution patterns typical of Boston aerosol. In a simulation study, we showed that the health effects estimates obtained using the tracer approach are attenuated, both in small- and large-sample cases, which was expected having framed the problem from a measurement error perspective. Our results suggest that the common marker for traffic particles, BC, has a relatively large degree of error associated with it, which may reflect a regional component of BC in addition to local traffic in the Boston area; as such the tracer approach may underestimate the effect associated with motor vehicles. The 2-stage approach is similarly susceptible to bias, although only in the case of small samples. For large samples, the 2-stage estimates appear unbiased as we would expect of a regression calibration.

As an alternative to the tracer and 2-stage approaches, we proposed a SEM to account for the uncertainty associated with latent source contributions, along with a Bayesian approach to model fitting. This approach leverages exposure information from previous related concentrator studies. Simulations suggest that the proposed informative Bayesian SEM is effective in eliminating bias in estimated source-specific health effects estimates, even when the number of exposures is limited. We demonstrated the flexibility of the Bayesian approach to accommodate complex study designs and nonnormality in our analysis of the study of Wellenius *and others* (2003). As an added advantage, the informative Bayesian SEM may be implemented in freely available software.

Our findings in this paper have implications for the design of future PM concentrator studies. The results demonstrate the benefits of using exposure data from existing, relevant exposure studies where possible. However, not all studies have the benefit of such prior knowledge. When historical data are not available, investigators should maximize the number of unique exposure days, subject to cost constraints. A large number of exposure days will allow one to use 2-stage models that employ source apportionment techniques to address errors in the tracer characterization of the pollution sources. Unfortunately, because even one run of a concentrator exposure can be costly, such “many-exposure” designs may be prohibitive. In cases where both historic and current exposure information is limited, the tracer approach is preferable to the 2-stage analysis, since the latter may be unstable and experience convergence problems; however, one should take care to choose “good” tracers that minimize the error associated with these surrogates. In these settings, the tracer approach is still useful in screening for source-specific health effects but is likely to yield attenuated estimates of effects.

The purpose of this article was to assess the performance of methods for source characterization in concentrator studies. We therefore focused on a specific form of a factor analysis model in the structural equations framework and have not addressed all the interesting modeling issues that arise in the development of a “good” receptor model. For instance, although the estimation of the number of sources can often be challenging (Park *and others*, 2002), we assumed that we have good prior knowledge on the number of major pollution sources in the Boston area. This assumption is probably reasonable in our setting as the pollution mixture in this area has been studied for almost 2 decades (Oh *and others*, 1997). We note that existing exposure studies suggest that there exists an additional particle source in the Boston aerosol comprised of sodium and chloride, often referred to as sea salt. However, from a regulatory perspective this exposure is not of primary importance and hence of less interest in PM health studies. Second, the majority of the source apportionment literature for exposure assessment of PM (Park *and others*, 2001) assumes that the sources of exposure are independent. To maintain consistency with existing methodology, we specified independent priors on the latent source contributions in the implementation of our Bayesian SEM. However, while source-specific exposures are assumed to be independent *a priori*, the Bayesian approach uses data to update the priors, thus allowing for correlation in the posterior distributions of the

source contributions. Given that we might expect source contributions to be correlated, at least in part due to meteorologic conditions, this is an appealing feature of our approach. Furthermore, the methods proposed in this paper extend naturally to account for systematic factors, like meteorology, that are likely to affect pollution levels from different sources similarly by allowing the means of the unobserved sources to depend on covariates.

ACKNOWLEDGMENTS

This research was supported by the National Institute of Environmental Health Sciences (NIEHS) grant ES07142 (Margaret C. Nikolov), American Chemistry Council grant 2843 and NIEHS grant ES012044 (Brent A. Coull), and National Institutes of Health grant ES012972 (John J. Godleski). The authors thank a referee for insightful comments that improved the manuscript. *Conflict of Interest*: None declared.

REFERENCES

- BANDEEN-ROCHE, K. (1994). Resolution of additive mixtures into source components and contributions: a compositional approach. *Journal of the American Statistical Association* **89**, 1450–1458.
- BATALHA, J. R. F., SALDIVA, P. H. N., CLARKE, R. W., COULL, B. A., STEARNS, R. C., LAWRENCE, J., MURTHY, G. G. K., KOUTRAKIS, P. AND GODLESKI, J. J. (2002). Concentrated ambient air particles induce vasoconstriction of small pulmonary arteries in rats. *Environmental Health Perspectives* **110**, 1191–1197.
- BOLLEN, K. A. (1989). *Structural Equations with Latent Variables*. New York: Wiley.
- CARROLL, R. J., RUPPERT, D. AND STEFANSKI, L. A. (1995). *Measurement Error in Nonlinear Models*. London: Chapman & Hall.
- CLARKE, R. W., COULL, B. A., REINISCH, U., CATALANO, P., KILLINGSWORTH, C. R., KOUTRAKIS, P., KAVOURAS, I., MURTHY, G. G. K., LAWRENCE, J., LOVETT, E. and others (2000). Inhaled concentrated ambient particles are associated with hematologic and bronchoalveolar lavage changes in canines. *Environmental Health Perspectives* **12**, 1179–1187.
- DOCKERY, D. W., POPE, C. A., XU, X., SPENGLER, J. D., WARE, J. H., FAY, M. E., FERRIS, B. G. AND SPEIZER, F. E. (1993). An association between air pollution and mortality in six U.S. cities. *New England Journal of Medicine* **329**, 1753–1759.
- DOMINICI, F., DANIELS, M., ZEGER, S. L. AND SAMET, J. M. (2002). Air pollution and mortality: estimating regional and national dose-response relationships. *Journal of the American Statistical Association* **97**, 100–111.
- DOMINICI, F., SHEPPARD, L. AND CLYDE, M. (2003). Health effects of air pollution: a statistical review. *International Statistical Review* **71**, 243–276.
- GODLESKI, J. J., CLARKE, R. W., COULL, B. A., SALDIVA, P. H. N., JIANG, N.-F., LAWRENCE, J. AND KOUTRAKIS, P. (2002). Composition of inhaled urban air particles determines acute pulmonary responses. *Annals of Occupational Hygiene* **46**, 419–424.
- GODLESKI, J. J., VERRIER, R. L., KOUTRAKIS, P., CATALANO, P., COULL, B. A., REINISCH, U., LOVETT, E. G., LAWRENCE, J., MURTHY, G. G. K., WOLFSON, J. M., CLARKE, R. W. AND NEARING, B. D. (2000). Mechanisms of morbidity and mortality from exposure to ambient air particulate. *Health Effects Institute Research Report* **91**, 1–103.
- HOPKE, P. K. (2003). Recent developments in receptor modeling. *Journal of Chemometrics* **17**, 255–265.
- LADEN, F., NEAS, L. M., DOCKERY, D. W. AND SCHWARTZ, J. (2000). Association of fine particulate matter from different sources with daily mortality in six U.S. cities. *Environmental Health Perspectives* **108**, 941–947.
- LAWRENCE, J., WOLFSON, J. M., FERGUSON, S., KOUTRAKIS, P. AND GODLESKI, J. J. (2004). Performance stability of the Harvard ambient particle concentrator. *Aerosol Science and Technology* **38**, 219–227.

- MUTHEN, L. K. AND MUTHEN, B. O. (1998). *Mplus User's Guide*, 3rd edition. Los Angeles, CA: Muthen & Muthen.
- NIKOLOV, M. C., COULL, B. A., CATALANO, P. J. AND GODLESKI, J. J. (2006). An informative Bayesian structural equation model to assess source-specific *s* of air pollution. *Harvard University Biostatistics Working Paper Series 46*. <http://www.bepress.com/harvardbiostat/paper46>. Accessed 25 October 2006.
- OH, J. A., SUH, H. H., LAWRENCE, J. E., ALLEN, G. A. AND KOUTRAKIS, P. (1997). Characterization of particulate mass concentrations in South Boston, MA. *Proceedings of AWMA/EPA Symposium on "Measurement of Toxic and Related Air Pollutants"*, April 29-May 1, 1997, Research Triangle Park, NC. Pittsburgh, PA: AWMA publication number VIP-74, pp. 397–407.
- PAATERO, P. AND TAPPER, U. (1994). Positive matrix factorization—a nonnegative factor model with optimal utilization of error-estimates of data values. *Environmetrics* **5**, 111–126.
- PARK, E. S., GUTTORP, P. AND HENRY, R. C. (2001). Multivariate receptor modeling for temporally correlated data by using MCMC. *Journal of the American Statistical Association* **96**, 1171–1183.
- PARK, E. S., SPIEGELMAN, C. H. AND HENRY, R. C. (2002). Bilinear estimation of pollution source profiles and amounts by using multivariate receptor models. *Environmetrics* **13**, 775–798.
- R DEVELOPMENT CORE TEAM (2003). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. <http://www.R-project.org>. Accessed 21 February 2006.
- SIOUTAS, C., KOUTRAKIS, P. AND BURTON, R. M. (1995). A technique to expose animals to concentrated fine ambient aerosols. *Environmental Health Perspectives* **103**, 172–177.
- SPIEGELHALTER, D., THOMAS, A. AND BEST, N. (2000). *WinBUGS Version 1.3. User's Manual*. Cambridge: MRC Biostatistics Unit, Institute of Public Health. <http://www.math.helsinki.fi/openbugs/data/Docu/WinBugs%20Manual.html>.
- TSIATIS, A. A., DE GRUTTOLA, V. AND WULFSOHN, M. S. (1995). Modeling the relationship of survival to longitudinal data measured with error; applications to survival and CD4 counts in patients with AIDS. *Journal of the American Statistical Association* **90**, 27–37.
- WELLENIUS, G. A., COULL, B. A., GODLESKI, J. J., KOUTRAKIS, P., OKABE, K., SAVAGE, S. T., LAWRENCE, J. E., MURTHY, G. G. K. AND VERRIER, R. L. (2003). Inhalation of concentrated ambient air particles exacerbates myocardial ischemia in conscious dogs. *Environmental Health Perspectives* **111**, 402–408.

[Received August 25, 2005; first revision May 1, 2006; second revision September 13, 2006;
accepted for publication October 3, 2006]