Journal of Exposure Analysis and Environmental Epidemiology (2005) 15, 366–376
© 2005 Nature Publishing Group All rights reserved 1053-4245/05/$30.00
www.nature.com/jea

Exposure and measurement contributions to estimates of acute air pollution effects

LIANNE SHEPPARD,[a,b] JAMES C. SLAUGHTER,[b] JONATHAN SCHILDCROUT,[a] L.-J. SALLY LIU[b] AND THOMAS LUMLEY[a]

[a]*Department of Biostatistics, University of Washington, Seattle, Washington, USA*
[b]*Department of Environmental and Occupational Health Sciences, University of Washington, Seattle, Washington, USA*

Air pollution health effect studies are intended to estimate the effect of a pollutant on a health outcome. The definition of this effect depends upon the study design, disease model parameterization, and the type of analysis. Further limitations are imposed by the nature of exposure and our ability to measure it. We define a plausible exposure model for air pollutants that are relatively nonreactive and discuss how exposure varies. We discuss plausible disease models and show how their parameterizations are affected by different exposure partitions and by different study designs. We then discuss a measurement model conditional on ambient concentrations and incorporate this into the disease model. We use simulation studies to show the impact of a range of exposure model assumptions on estimation of the health effect in the ecologic time series design. This design only uses information from the time-varying ambient source exposure. When ambient and nonambient sources are independent, exposure variation due to nonambient source exposures behaves like Berkson measurement error and does not bias the effect estimates. Variation in the population attenuation of ambient concentrations over time does bias the estimates with the bias being either positive or negative depending upon the association of this parameter with ambient pollution. It is not realistic to substitute measured average personal exposures into time series studies because so much of the variation in personal exposures comes from nonambient sources that do not contribute information in the time series design. We conclude that general statements about the implications of measurement error need to be conditioned on the health effect study design and the health effect parameter to be estimated.
Journal of Exposure Analysis and Environmental Epidemiology (2005) **15,** 366–376. doi:10.1038/sj.jea.7500413
Published online 15 December 2004

Keywords: *air pollution, acute effects, time series studies, panel studies, measurement error*

Introduction

The goal of many air pollution epidemiology studies is to estimate the effect of air pollution on health. One challenging aspect of this task is that the meaning of the health effect parameter depends upon the study design, the disease model (i.e., the assumed relationship between air pollution and exposure), the approach to analysis, the alignment of exposure measurements with actual exposure, as well as many other factors. Thus, there are different parameters that all can be called health effects. In addition, conditional on the parameters, in any particular study there are impacts on estimates due to confounding, measurement error, and the interplay of the data with the study design and analysis approach. The goal of this paper is twofold: to point out some of the factors that affect what parameter is being estimated based on the design of a health effect study, and to give new insights into the bias and variance of the estimates actually achieved in a specific health effect study. For both goals, we focus on short-term time-varying personal exposure and its effect on individual health measures. For the first goal, we discuss the parameterization of various disease models in the context of the exposure distribution and exposure measurements. To achieve the second goal, we need to add plausible exposure and measurement models; we use data from Seattle to guide parameter choice in the hypothesized models. We also restrict our attention to the time series study design, but apply it in the context of an underlying individual-level disease model for time-varying exposure that is generating the data. Through designed simulation studies, we show the sensitivity of the health effect estimates to changes in the exposure distribution and its measurement.

A disease model takes the general form

$$\mathrm{E}(Y) = f(X, \beta, Z, \gamma) \tag{1}$$

where Y is a disease outcome, X is the exposure, β is the exposure effect parameter of interest, Z is a vector of other risk factors or confounders, γ is Z's associated vector of parameters, and $f(\cdot)$ describes their functional relationship. The definition of β depends upon the disease model since the functional relationship, $f(\cdot)$, and which Z's are included will

1. Address all correspondence to: Lianne Sheppard, Box 357232, Department of Biostatistics, University of Washington, Seattle, WA 98195-7232, USA. Tel.: +206 616 2722. Fax: +206 616 2724.
E-mail: sheppard@u.washington.edu
Received 18 June 2003; accepted 19 July 2004; published online 15 December 2004

change its interpretation (e.g., as an additive *versus* relative risk, or as an exposure effect when other risk factors are held constant). It is less well appreciated that the parameter β can have interpretations that also depend on the exposure distribution, analysis approach, and study design. We discuss some examples in the next section.

Time series studies have been the bread and butter of air pollution epidemiology. In spite of their widespread use, there has been little work to evaluate what parameters these studies are actually estimating, or to assess the ability of these studies to estimate the target parameters in light of the available exposure data. Measured exposures in air pollution time series studies are ambient concentrations from fixed-site ambient monitors. These are treated as fixed and known in the analysis. As we discuss later, the disconnection between the exposure measurements and the target exposure (e.g., total personal exposure) is an important source of bias. With a few exceptions, interpretation of the health effects estimated in the time series study literature has not explicitly addressed the disconnection between ambient concentration and personal or population exposure (Dominici et al., 2000; Sheppard and Damian, 2000).

Models for estimating exposure effects in the presence of exposure measurement error provide a framework for clarifying what parameters are being estimated in air pollution studies such as time series studies. The general approach to measurement error modeling is to specify three models: the disease model, the exposure model, and the measurement model (Clayton, 1991; Gilks and Richardson, 1992). The *disease model*, for example, (1), relates a disease outcome to specific exposures. Standard disease model analyses condition on the exposure X being fixed and known. Typically, the exposures X are not known or measured; these are true underlying or latent exposures. The *exposure model* specifies the distribution of latent exposures in the population. The *measurement model* relates the measured exposures to the latent exposures. A necessary assumption for this formulation is that no additional information about the disease is contained in the exposure measurements, above and beyond that contained in the underlying exposures. While complex, there are measurement error modeling techniques that can be applied to estimate β in the presence of exposure measurement error (Carroll et al., 1995).

Measurement error is part of a priority research topic for airborne particulate matter identified by the National Research Council Committee on Research Priorities for Airborne Particulate Matter (1998). The committee defines measurement error as "the difference between actual exposures and measured ambient air concentrations" (National Research Council Committee on Research Priorities for Airborne Particulate Matter, 2001). They state that "the three sources of measurement error are instrument error (the accuracy and precision of the monitoring instrument), error resulting from the nonrepresentativeness of a monitoring site (reflected by the spatial variability of the pollutant measured), and differences between the average personal exposure to a pollutant and the monitored concentration (influenced by microenvironmental exposures)." This follows from work on measurement error in air pollution time series studies (Dominici et al., 2000; Zeger et al., 2000). Although these papers are specific to the time series design, the NRC report implies that the three sources of measurement error are important regardless of study design. We believe that the context of the study — its design, analysis, the outcome, the available exposure data, and the latent exposure distribution — will combine to magnify or diminish each of these measurement error sources.

We suggest that it is valuable to use the measurement error modeling framework to gain understanding about disease model parameters in specific contexts. Then under the assumptions required by the measurement error modeling, one can assess the impact measurement error and exposure distribution assumptions have within the context of a particular study design. After incorporating factors such as the effect size and analysis approach, one can determine the consequences of ignoring the measurement error problem. This is how we approach the second aim of this paper where we assess the effect of exposure variation and measurement in the time series design.

Role of exposure, disease model and study design

Personal Exposure Model

We restrict our attention to a (relatively) nonreactive pollutant (such as particulate matter (PM) or carbon monoxide (CO)) and assume that total personal exposure to that pollutant is the relevant exposure for any health model. We concentrate on relatively nonreactive pollutants because reactive pollutants (e.g., ozone) will be more sensitive to the complex interactions that individuals have with their environments and less likely to satisfy approximately the assumptions of our proposed model. Our attention on total personal exposure will give insight into the interpretation of health effect parameters and the implications of study design. Most epidemiologic studies rely on ambient (outdoor) concentration measurements from central site monitors. Thus, by virtue of the available "exposure" measurement data, these epidemiologic studies target ambient exposures. Ambient source exposures are those derived from the outdoor shared environment. Ambient exposures are also interesting because policies regulate ambient source pollutants differently than nonambient source pollutants. (Note that for particulate matter, chemical composition and dominant sources may be very different for ambient and nonambient source exposures.)

A model for total personal exposure for individual i at time t, X_{it}^{P}, can be divided into two dominant sources, ambient

(X_{it}^{A}) and nonambient (X_{it}^{N}) source exposures:

$$X_{it}^{P} = X_{it}^{N} + X_{it}^{A} = X_{it}^{N} + \alpha_{it} X_{it}^{C}. \quad (2)$$

Nonambient sources are derived from indoor, personal, or local environments. Note that X_{it}^{C} is ambient concentration at individual i's spatial location (and only equals ambient source exposure if the individual spends all their time outdoors), and α_{it} is the degree of attenuation from the ambient concentration (Ott et al., 2000; Sheppard and Damian, 2000; Wilson et al., 2000). We further assume that X_{it}^{N} is independent of X_{it}^{A} and α_{it} is independent of X_{it}^{C}. Our model differs from an exposure model that weights the nonambient source component by $(1-\alpha_{it})$ (Dominici et al., 2000; Zeger et al., 2000), which does not coincide with the mass balance equation interpretation for α_{it} (see below).

With respect to study design, we believe that dividing exposure into two dominant sources is likely to be the most important partitioning of exposure. More elaborate partitions that are based on Duan's original microenvironmental concept (Duan, 1991) include the probability- and population-based exposure model (Burke et al., 2001) and the computational personal exposure model (Zidek et al., 2000, 2003). Such microenvironmental models do not distinguish explicitly ambient and nonambient contributions to personal exposure.

A model for the nonambient source component is

$$X_{it}^{N} = (\beta_0 + b_i) + \varepsilon_{it} \quad (3)$$

where β_0 represents an average nonambient concentration across all individuals and times in a population, b_i is a subject-specific intercept representing the deviation of an individual's usual exposure from the population average, and ε_{it} represents within-person deviations from usual exposure over time. Typically, one would assume ε_{it} to have a normal distribution with mean zero and variance σ_{it}^2, b_i to have either a normal or lognormal distribution, and b_i and ε_{it} to be independent. This model can be extended to incorporate person- or time-specific attributes into the intercept (i.e. exposure predictors such as type of home heating, cooking behavior), dependence over time, or other distributional structures.

The ambient source component of personal exposure has two parts. The first is an attenuation parameter to adjust the ambient concentration to an individual's actual ambient exposure. The attenuation parameter α_{it} depends upon a person's behavior, their microenvironment, and the qualities of the pollutant. The chemical mass balance equation gives the underlying structure for modeling α_{it} (Koutriakis et al., 1992). The key features making up the equation are the penetration (P), deposition (k) and decay (a) rates of the pollutant along with human behavior, specifically the fraction of time an individual spends outdoors (o_{it}). A plausible model for α_{it} is therefore

$$\alpha_{it} = o_{it} + (1 - o_{it}) \frac{P}{a+k}$$

(Mage et al., 1999). For particulate matter, this formulation ignores exposure due to resuspension of ambient PM in the local environment.

The second part of the ambient source component of personal exposure is the ambient concentration, X_{it}^{C}. This is known to vary over time because of changes in source generation, weather, and season. Ambient concentrations also vary over space as a function of geography (specifically topography that provides barriers or encourages drainage flow), weather conditions (that encourage or inhibit environmental mixing), and local sources (both point and mobile sources). It is reasonable to expect variation over space and time in ambient concentrations to differ regionally and by pollutant.

The 'Exposure data and parameters for health studies' section summarizes $PM_{2.5}$ data from Seattle under this model.

Panel Study Disease Model

A traditional panel study enrolls individuals $i = 1, \ldots, N$ and observes them at times $t = 1, \ldots, T$. One possible health effects model for a continuous outcome is

$$\mathrm{E}(Y_{it} | X_{i(t-\ell)}^{P}, Z_{it}) = \gamma_0 + X_{i(t-\ell)}^{P}\beta + g(Z_{it}, \gamma_Z) \quad (4)$$

where ℓ indicates the lag of the air pollutant, Z_{it} are confounders, and (γ_0, β, γ_Z) are parameters to be estimated. We assume the pollutant effect β enters linearly in the model, but there is no such constraint on the confounder effects (i.e., $g(\cdot)$ is an arbitrary function). The linearity assumption is supported by dose–response modeling in time series studies (Dominici et al., 2002; Schwartz et al., 2002), although such studies may not have much power to detect departures from linearity. We also assume that there is only a single lag for exposure. Multiple lags or distributed lag alternatives could be considered as well.

The comparable models for binary outcomes are

$$\mathrm{E}(Y_{it} | X_{i(t-\ell)}^{P}, Z_{it}) = exp\{\gamma_0 + X_{i(t-\ell)}^{P}\beta + g(Z_{it}, \gamma_Z)\} \quad (5)$$

for a rare event, and

$$\mathrm{E}(Y_{it} | X_{i(t-\ell)}^{P}, Z_{it}) = exp\{\gamma_0 + X_{i(t-\ell)}^{P}\beta + g(Z_{it}, \gamma_Z)\} / (1 + exp\{\gamma_0 + X_{i(t-\ell)}^{P}\beta + g(Z_{it}, \gamma_Z)\}) \quad (6)$$

for a common outcome.

Types of Effects

The parameterization of the exposure effect described above can be expanded. Most commonly in panel studies, we recommend summarizing the average exposure for each individual, $\overline{X_i^P} = \Sigma_{t=1}^{T} X_{it}^{P}/T$ and allowing this term to have a separate parameter in the model. For instance, (5) becomes

$$\mathrm{E}(Y_{it} | X_{i(t-\ell)}^{P}, Z_{it}) = exp\{\gamma_0 + (X_{i(t-\ell)}^{P} - \overline{X_i^P})\beta_{W_p} + \overline{X_i^P}\beta_{B_p} + g(Z_{it}, \gamma_Z)\} \quad (7)$$

where β_{W_p} is the effect of exposure as it varies within a person, and β_{B_p} is the effect of average exposure as it varies between people. The assumption that β_{B_p} equals β_{W_p} is the assumption of no *contextual effects* (Firebaugh, 1978; Sheppard, 2003), which is often reasonable for environmental exposures (Sheppard, 2002). However, even when the underlying parameters are identical, their estimates can be very different due to differential effects of uncontrolled confounding and exposure measurement error (Sheppard, 2003).

This is an example where study design and exposure data can restrict the parameters that can be estimated. In a traditional panel study, all individuals are observed over the same time period. Furthermore, because all subjects are observed simultaneously and when only ambient concentration measurements are used for the exposure (i.e. substitute X_t^C for X_{it}^P), $\overline{X_i^C}$ will not vary across individuals. Without variation in $\overline{X_i^C}$, β_{B_p} cannot be estimated separately, but rather becomes absorbed into the intercept. Alternatively, when the time period of observation varies for each member of the panel, there is variation in $\overline{X_i^C}$ with which to estimate β_{B_p}. However, all the variation is due to the time period under which that person was studied and therefore bias due to seasonal confounding will likely be present. We have argued for estimating but not reporting β_{B_p} due to the likelihood of residual confounding. (see, e.g., Yu et al., 2000).

Another partition separates population temporal variation from the rest of the exposure by defining $\overline{X_t^P} = \Sigma_{i=1}^N X_{it}^P/N$. The associated model is given by

$$\begin{aligned} E(Y_{it}|X_{i(t-\ell)}^P, Z_{it}) = & exp\{\gamma_0 + (X_{i(t-\ell)}^P - \overline{X_{(t-\ell)}^P})\beta_{W_t} \\ & + \overline{X_{(t-\ell)}^P}\beta_{B_t} + g(Z_{it}, \gamma_Z)\} \end{aligned} \quad (8)$$

where β_{W_t} is the effect of a person's deviation in exposure from the shared temporally varying population exposure, while β_{B_t} is the effect of the shared population exposure as it varies over time. In a panel study the population is defined as all members of the panel.

This is another example where the exposure data will have a large impact on what parameters are possible to estimate. Again, substituting a single ambient exposure concentration measurement for total personal exposure in the model, the difference $(X_{i(t-\ell)}^C - \overline{X_{(t-\ell)}^C})$ will always be 0, resulting in no information with which to estimate β_{W_t}. While the assumption of no contextual effects implies $\beta = \beta_{W_t} = \beta_{B_t} = \beta_{W_p} = \beta_{B_p}$, the presence of exposure measurement error and residual confounding will make it likely that the estimates of all these parameters differ.

Role of Analysis

Fixed and Mixed Models Analysis approach can play a crucial role in what information enters into the parameter estimates. For instance, (5) can be expanded to include an individual susceptibility term h_i:

$$\begin{aligned} E(Y_{it}|X_{i(t-\ell)}^P, Z_{it}) = & h_i exp\{\gamma_0 + X_{i(t-\ell)}^P\beta \\ & + g(Z_{it}, \gamma_Z)\}. \end{aligned} \quad (9)$$

Suppose this is the model to be estimated while the true model is

$$\begin{aligned} E(Y_{it}|X_{i(t-\ell)}^P, Z_{it}) = & h_i exp\{\gamma_0 + (X_{i(t-\ell)}^P - \overline{X_i^P})\beta_{W_p} \\ & + \overline{X_i^P}\beta_{B_p} + g(Z_{it}, \gamma_Z)\} \end{aligned} \quad (10)$$

Analysis of (9) can treat each h_i as separate fixed effects (stratified analysis). Alternatively, we can allow h_i to have a distribution with mean 1 and variance σ_h^2 (random effects analysis). When $\sigma_h^2 = 0$, all individuals are constrained to have the same intercept and we have (5). Sheppard and Prentice (1995) discuss how the model for h_i affects the estimate of β. In a model that does not partition effects (e.g. (9)) and in the absence of any confounding, the parameter estimated in the stratified analysis is purely β_{W_p}, while for random effects analyses it is a weighted average of β_{W_p} and β_{B_p} with relative weights determined by σ_h^2 and the between *versus* within group variation in exposure (Piantadosi et al., 1988; Sheppard and Prentice, 1995).

Aggregation Aggregation can be done in the exposure alone, in the outcome and associated disease model, or both. It can be viewed as an aspect of study design or the analysis approach. Sheppard (2003) discusses the role of aggregation in exposure *versus* the disease model and the implications for estimation of exposure effects in the typical setting where individuals are nested within areas. For acute effect studies, disease models (4), (5) or (6) can be aggregated by summing or averaging over all individuals in the population on each day, for example, considering a model for $E(Y_t)$ where $Y_t = \Sigma_i Y_{it}$ or $Y_t = \Sigma_i Y_{it}/N$. This is the conceptual basis for time series studies. In an *aggregate model*, the model for $E(Y_{it})$ is specified at the individual level and then aggregated. It is obtained from directly aggregating disease models over individuals in the population. Aggregation of the panel study model is different from other aggregate study models because individuals are not nested within time periods. Instead, individuals are crossed with time and thus have their exposure represented within each of the aggregates.

Aggregation of the disease model is another way to focus how the information in the exposure is used in estimating the health effect parameter. In this section, we use (5) because a model for rare events is the appropriate one for eventual application to air pollution time series studies. It also allows us to show the issues that arise when the link function is not linear. Linear models (e.g., (4)) are simpler to aggregate while logistic models (e.g., (6)) pose greater challenges after aggregation (Salway and Wakefield, in press). Using (5),

the aggregate model is

$$\mathrm{E}(Y_t) = \sum_i exp\{\gamma_0 + X_{it}^{\mathrm{P}}\beta + g(Z_{it}, \gamma_Z)\}. \quad (11)$$

Alternatively, aggregation can be confined to just the exposure variable. A *semi-individual model* is still an individual-level model, but in this case the exposure represents a population-level exposure. For instance, substituting the within-population, temporally averaged component of exposure in (5) yields the semi-individual model

$$\mathrm{E}(Y_{it}|\overline{X_{(t-\ell)}^{\mathrm{P}}}, Z_{it}) = exp\{\gamma_0 + \overline{X_{(t-\ell)}^{\mathrm{P}}}\beta + g(Z_{it}, \gamma_Z)\}. \quad (12)$$

When the disease model does not have a linear link function, aggregation of the exposure without regard for the disease model can be a source of bias called *specification bias* (Sheppard, 2003). The bias comes from not using within-day exposure variation in estimating β and it occurs because $\mathrm{E}(\exp\{X\beta\}) \neq \exp\{\mathrm{E}(X)\beta\})$. Common wisdom suggests that there is no bias because the measurement $\overline{X_{(t-\ell)}^{\mathrm{P}}}$ has Berkson error structure (Kunzli and Tager, 1997; Zeger et al., 2000). However, in the case of a nonlinear disease model where the variance of the unmeasured part of the exposure is correlated with the mean (as would be common for a lognormally distributed exposure), there can be specification bias due to using an average of the exposure in the disease model (Sheppard, 2003). Specification bias will be small when β is small; this is often true for air pollution health effect studies.

Unlike the semi-individual study, an ecological study design uses population-level exposures and confounders. It also does the analysis on (a version of) the aggregated disease model. Once all variation at the individual level has been removed from the model predictors, the aggregation does not change the disease model. Thus modifying the confounders to only be included at the population level and then aggregating (12) gives an *ecological time series model*:

$$\mathrm{E}(Y_t|\overline{X_{(t-\ell)}^{\mathrm{P}}}, Z_t) = exp\{\gamma_0^e + \overline{X_{(t-\ell)}^{\mathrm{P}}}\beta^e + g(Z_t, \gamma_Z^e)\}. \quad (13)$$

Here, $\overline{X_{(t-\ell)}^{\mathrm{P}}}$ is the average personal exposure on each day for the entire population. In practice, $\hat{X}_{(t-\ell)}^{\mathrm{C}}$, an exposure measurement from an ambient monitor, is substituted for $\overline{X_{(t-\ell)}^{\mathrm{P}}}$ in this model. The time-varying confounders most often included are smooth functions of time (to control for season, trends, and epidemics), day of week indicators, and smooth functions of weather variables, specifically temperature, dew point, and relative humidity. While these are not population average confounders, they are believed to capture the risk factors that are correlated with pollution exposure.

Specification of the model at the individual level and then introducing aggregation allows us to highlight several research directions, including: implications of substituting the ecologic for the aggregate model, and what is the impact of using ambient monitor measurements for total personal exposure. We explore these more in the following section and consider them from a practical point of view in the penultimate section.

Incorporating a measurement error model

Under the assumption that the underlying exposure is X_{it}^{P}, and only time-varying ambient concentration X_{it}^{C} is measured, analysis can still be done by specifying a model for the distribution of $X_{it}^{\mathrm{P}}|X_t^{\mathrm{C}}$. Under the disease model (5), the model of interest is

$$\mathrm{E}(exp\{\gamma_0 + X_{i(t-\ell)}^{\mathrm{P}}\beta + g(Z_{it}, \gamma_Z)\}|X_{(t-\ell)}^{\mathrm{C}})$$

It is necessary to specify $g(\cdot)$ and a model for the joint exposure measurement and confounding distribution to simplify this expectation. In order to concentrate on the effect of exposure measurement error in isolation of its effects on confounding, we consider the simpler model:

$$\mathrm{E}(exp\{\gamma_0 + X_{i(t-\ell)}^{\mathrm{P}}\beta\}|X_{(t-\ell)}^{\mathrm{C}})$$

Since (5) is not linear, it is necessary to specify $\mathrm{E}(X_{it}^{\mathrm{P}}|X_t^{\mathrm{C}})$ and $\mathrm{var}(X_{it}^{\mathrm{P}}|X_t^{\mathrm{C}})$. These are needed for an approximation to the expectation of a moment generating function when $X_{it}^{\mathrm{P}}|X_t^{\mathrm{C}}$ is not normally distributed. (The result is exact when it is normally distributed.) It simplifies to

$$exp\{\gamma_0 + \mu_{X^{\mathrm{P}}|X_t^{\mathrm{C}}}\beta + \sigma^2_{X^{\mathrm{P}}|X_t^{\mathrm{C}}}\beta^2/2\} \quad (14)$$

Under the assumptions for the exposure model in the first section and for $\mathrm{E}(\alpha_{it}) = \alpha$,

$$\begin{aligned}\mu_{X^{\mathrm{P}}|X_t^{\mathrm{C}}} &= \beta_0 + \alpha\mathrm{E}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}})\\ \sigma^2_{X^{\mathrm{P}}|X_t^{\mathrm{C}}} &= \mathrm{var}(b) + \mathrm{var}(\varepsilon)\\ &\quad + \mathrm{var}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}})[\alpha^2 + \mathrm{var}(\alpha_{it})]\\ &\quad + \mathrm{E}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}})^2\mathrm{var}(\alpha_{it}).\end{aligned}$$

Determining $\mathrm{E}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}})$, the components of $\sigma^2_{X^{\mathrm{P}}|X_t^{\mathrm{C}}}$, and their correlation is the first step in assessing the impact of merely substituting X_t^{C} for X_{it}^{P} in a health effects model.

These results can be combined with model (11) from the 'Role of exposure, disease model and study design' section to show how spatial variation in the ambient concentration in air pollution time series studies affects the parameter estimates. When there is pure spatial variation in ambient concentration, $\mathrm{E}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}})$ will depend on space only through a spatially varying intercept, for example, $\mathrm{E}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}}) = \delta_{0i} + \delta_1 X_t^{\mathrm{C}}$. The same would hold for the conditional variance, for example, $\mathrm{var}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}}) = \delta_{0i}^* + \delta_1^*\mathrm{var}(X_t^{\mathrm{C}})$. Under the additional restrictions that $\mathrm{var}(b) = \tau_1$, $\mathrm{var}(\varepsilon) = \tau_2$, $\mathrm{var}(\alpha_{it}) = \tau_3$ and there are no confounders, aggregating the disease model (14) as for (11) yields

$$\begin{aligned}\mathrm{E}(Y_t) = \sum_i exp\{&\gamma_0 + (\beta_0 + \alpha(\delta_{0i} + \delta_1 X_t^{\mathrm{C}}))\beta\\ &+ (\tau_1 + \tau_2 + (\delta_{0i}^* + \delta_1^*\mathrm{var}(X_t^{\mathrm{C}}))[\alpha^2 + \tau_3]\\ &+ (\delta_{0i}^2 + 2\delta_{0i}\delta_1 X_t^{\mathrm{C}} + \delta_1^2(X_t^{\mathrm{C}})^2)\tau_3)\beta^2/2\}.\end{aligned}$$

Collecting terms that do not depend on time into γ_{0i}^* and simplifying, we obtain

$$\begin{aligned}\mathrm{E}(Y_t) &= \sum_i exp\{\gamma_{0i}^* + \alpha\delta_1 X_t^{\mathrm{C}}\beta \\ &+ ((\delta_1^* \mathrm{var}(X_t^{\mathrm{C}})[\alpha^2 + \tau_3] \\ &+ (2\delta_{0i}\delta_1 X_t^{\mathrm{C}} + \delta_1^2 (X_t^{\mathrm{C}})^2)\tau_3)\beta^2/2\} \\ &= exp\{\alpha\delta_1 X_t^{\mathrm{C}}\beta + [\delta_1^2 (X_t^{\mathrm{C}})^2\tau_3 \\ &+ [\alpha^2 + \tau_3]\delta_1^* \mathrm{var}(X_t^{\mathrm{C}})]\beta^2/2\} \\ &\times \sum_i exp\{\gamma_{0i}^* + \delta_{0i}\delta_1 X_t^{\mathrm{C}}\beta^2\}.\end{aligned}$$

The last term aggregates a function of concentration over space. Time and space are separable here, so it is only the nonlinearity in the function exp(·) that induces any temporal variation in the spatial variance. Thus, the spatial variability only induces specification bias into the disease model. However, because β^2 will be quite small in air pollution applications, there will be limited specification bias. Thus there will be little effect of the spatial variability of concentration on the disease model. Note that the parameter β in this aggregated induced model differs from the parameter β^ein the related time series model

$$\mathrm{E}(Y_t) = exp\{\gamma_0^{\mathrm{e}} + X_t^{\mathrm{C}}\beta^{\mathrm{e}}\}.$$

In contrast, when there is spatiotemporal variation in the exposure, models for $\mathrm{E}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}})$ and $\mathrm{var}(X_{it}^{\mathrm{C}}|X_t^{\mathrm{C}})$ will vary over time by location. This exposure distribution will have a greater impact on the time series model. Many more terms in the disease model will be affected by the spatial variability. Then the practical question becomes how large is the spatiotemporal variation and thus how much will it affect the health effect estimate. If it is large, it may be possible to identify predictors such as weather that may explain much of this source of variation.

We now extend this analytic exercise into a practical one by simulating a variety of exposure and measurement scenarios. The following section sets the stage for the simulations discussed in the subsequent penultimate section by using data to suggest reasonable assumptions for the exposure and measurement distributions. In the penultimate section, we simulate data under an individual-level model for total personal exposure and its effect on health. We ignore this structure in the analysis and analyze the data using the time series study design. Our goal is to understand the implications in practice of the interplay between exposure distribution, measurement error distribution, and study design to health effect estimates from time series studies.

Exposure data and parameters for health studies

Example: Data on $PM_{2.5}$ in Seattle

We used information from a panel study being conducted in Seattle to estimate the heterogeneity of ambient concentration, the local ambient exposure, the attenuation of the ambient exposure, and variability in the nonambient exposure. All data were collected between October 1999 and June 2001. For the personal exposure model we used data on 133 elderly or children subjects residing in 117 locations who were monitored for exposure to particulate matter beginning in February 2000. Harvard personal environmental monitors (HPEMs) measured personal exposure to $PM_{2.5}$ and Harvard impactors (HIs) located outside of a subject's residence estimated local ambient concentration. A HI run concurrently at a central site estimated ambient exposure. Measurements were taken during 22 10-day sessions (average of seven subjects per session) with some subjects monitored twice or three times (Liu et al., 2002). For assessing spatial heterogeneity of ambient concentration we used nephelometers from 65 home sites over 24 10-day sessions. In addition, we included measurements from nephelometers located at two fixed locations throughout the entire period of panel study observation. All measurements for this analysis were for a study day (from 1600 hours of a day to 1600 hours of the following day; nephelometers measured continuously and readings were averaged). Liu et al. (2002) show daily average nephelometer light scattering data are highly correlated with gravimetric $PM_{2.5}$, so we converted nephelometer values to the same gravimetric scale used for the HPEMs and HIs (μg/m^3) using their regression model.

Estimates of Instrument Error

Sampler evaluation parameters for mass-based samplers such as HPEMs, HIs, and the Federal Reference Method (FRM) monitors include detection limit, size cut, and inaccuracy (bias and imprecision). Personal exposures are difficult to measure accurately due to the size and noise limitations of the samplers and pumps. Factors potentially influencing sampler performance include temperature, relative humidity, analysis (weighing) errors, sampler application (e.g., stationary *versus* personal), and sample handling. From the Seattle panel study we estimated the imprecision to be 1.2 μg/m^3 (or 8%) for HIs and 2.2 μg/m^3 (or 17%) for HPEMs. The bias based on comparisons with the FRM method is 1.1 μg/m^3 (or 10%) for HIs and 1.5 μg/m^3 (or 18%) for HPEMs (Liu et al., 2002).

Personal $PM_{2.5}$

We. have measurements for X_{it}^{P}, X_{it}^{C} and X_t^{C} in μg/m^3 from personal, home outdoor and central site monitors, respectively. We restrict our dataset to person-times with no missing measurements. We fit the random component superposition (RCS) model (Ott et al., 2000) to (2) under the strong assumption that $\alpha_{it} = \alpha$ and using X_t^{C} as the predictor. In this case we obtain $\alpha = 0.49$ (95% CI: 0.36, 0.62), and also using (3) to model X_{it}^{N}, $\beta_0 = 5.4$, $\mathrm{var}(b_i) = 20.7$, and $\mathrm{var}(\varepsilon_{it}) = 20.5$. RCS analysis using X_{it}^{C} as a predictor yielded nearly identical

results with a slightly smaller estimate of an estimate of var(b_i) (19.7). It is important to note that this analysis does not account for instrument error. Instrument error in X_{it}^{P} (estimated at 4.84) is incorporated into the estimate of var(ε_{it}) while instrument error in X_{it}^{C} and X_t^{C} has the effect of attenuating the estimate of α. Using method of moments approach and a classical measurement error model with measurement error variance of 1.44 (HI imprecision^2), the corrected estimate of α is 0.53 (95% CI: 0.40, 0.67). Further exploration of the data indicated that a lognormal assumption for the distribution of b_i was more plausible than normality, but that ε_{it} could reasonably be assumed to behave like a normally distributed variable. Using our data and the RCS model assumptions, we estimate that 10–12.5% of the variation in personal exposure in the panel study is due to ambient exposure (the higher estimate is after removing estimates of instrument error). Improved estimates of α may revise this percentage. Related work suggests the constant α assumption in the RCS model is too strong: α varies considerably across individuals and is higher during the nonheating rather than the heating season (Allen et al., 2003).

Ambient $PM_{2.5}$: Role of Spatial Heterogeneity

We fit an analysis of variance model to the daily average ambient nephelometer data from 67 locations (65 homes plus two central site locations) over 238 days of observation. Using a separate estimate for each date and location, we estimate that 87% of the variation in this dataset can be explained by time, 7% by spatial location, and 6% is unexplained. The low amount of unexplained variation suggests very little spatiotemporal variation in $PM_{2.5}$ in Seattle. The square root of the residual variation is 2.2 μg/m^3. Since we fit the sums of squares sequentially beginning with time and allow each date to be fit separately, our estimate that 87% of the variation is explained by time is optimistic. Alternatively, assuming the average of the two central sites is a good measurement of X_t^{C}, we find that X_t^{C} explains 84% of the variation in X_{it}^{C} across the 67 sites in a regression analysis. The estimate of residual error from that analysis is 3.1.

We use the regression of X_{it}^{C} on X_t^{C} (where X_t^{C} is measured by the average of two central sites) further to quantify $E(X_{it}^{C}|X_t^{C})$ and $\mathrm{var}(X_{it}^{C}|X_t^{C})$. Assessing the residuals from the regression analysis, we find that there is strong evidence of a mean-variance relationship in these data such that $\mathrm{var}(X_{it}^{C}|X_t^{C}) \doteq (0.037)\ E(X_{it}^{C}|X_t^{C})^2$. We estimate $E(X_{it}^{C}|X_t^{C})$ to be $0.9X_t^{C}$–1.6 μg/m^3.

Consequences for estimation in time series studies

The goal of this section is to determine the impact of a range of variants in the PM distribution on the ability to estimate the health effect parameter in an ecologic time series study. This is accomplished through a series of simulation studies. We consider the role of various distributional assumptions for the exposure under a specified model for the outcome. We do not directly address the effects of aggregation in our analysis. However, any bias due to aggregation will be manifested in the health effect parameter estimates we report.

Simulation Structure

For the foundation of these studies, we focus on the type of data that would generally be available for an ecologic time series analysis. Typically, these are routinely reported health events that are rare for individuals in the population. For feasibility reasons, we restrict the time period to three years ($T = 1000$ days) and the population size to $N = 100{,}000$ individuals. We scale the event rates so approximately 6 events are expected per day.

In the basic simulation structure we first simulate N individual exposures over T days according to the hypothesized exposure distribution. This gives us realized data from our exposure model. We then use these simulated exposures and a model for the outcome to simulate binary responses for each day. This gives us realized data from the individual-level outcome model. The outcomes are then aggregated by day. Then we estimate the health effect parameter β by fitting the ecologic time series model that most closely corresponds to the aggregate model we derived, given the realities of the available exposure data. We use some version of the exposure as the predictor. We consider the latent ambient concentration (known in simulations) as well as monitor and personal measurements. This whole process is repeated 500 times in order to allow us to summarize properties of $\hat{\beta}$. We address in sequence issues with the ambient *versus* nonambient source exposures, the distribution of the fraction of ambient concentration an individual is exposed to, and the effect of personal exposure variation (both spatial variation in ambient concentrations and use of personal exposure measurements in the time series model).

Contribution of Nonambient Source Exposure

For the first set of simulations we assume a constant attenuation, $\alpha = 0.5$, and that the ambient concentration is constant over all individuals in the population, that is, $X_{it}^{C} = X_t^{C}$. We used measured ambient values from a fixed site in Seattle for the period 1997–1999 for our known ambient source data. We use model (3) for the nonambient source exposure (X_{it}^{N}) with b_i and ε_{it} allowed to have either normal or lognormal distributions. The parameters σ_b^2, and σ_ε^2 vary across simulations, while the parameters β_0, μ_ε and μ_b are set so $E(X_{it}^{N}) = 39.25$ regardless of whether b_i or ε_{it} were normally or lognormally distributed. We set $\gamma_0 = ([\mathrm{var}(b_i) + \mathrm{var}(\varepsilon_{it})]\beta^2/2 - \beta_0 * \beta - 10.64)$ to induce a constant event rate across simulations.

Table 1 gives the results of these simulations averaged over 500 replications. The true value for β is $\ln(1.2)/10 = 0.0182$.

Table 1. Estimates of $\beta^e = \alpha\beta$ under independence of ambient and nonambient source exposures and a range of distributional assumptions for the nonambient source exposure.

b_i		ε_{it}					95% Coverage	
Distribution	var(b_t)	Distribution	var(ε_{it})	$\widehat{\alpha\beta}$[a]	$\sqrt{\text{var}(\widehat{\alpha\beta})}$[a]	$\sqrt{\widehat{\text{var}(\widehat{\alpha\beta})}}$[a]	$\alpha\beta$	β
—	0	Normal	107	9.3	2.9	3.0	0.93	0.13
—	0	Normal	161	8.9	3.0	3.0	0.95	0.11
—	0	Lognormal	74	8.8	2.6	2.7	0.95	0.05
—	0	Lognormal	134	9.0	2.5	2.6	0.96	0.05
Normal	49	Normal	68	9.0	2.9	3.0	0.97	0.12
Normal	49	Normal	102	9.0	3.0	3.0	0.95	0.11
Normal	74	Normal	68	9.2	3.0	3.0	0.95	0.15
Normal	74	Normal	102	9.0	2.9	3.0	0.96	0.12
Lognormal	46	Normal	68	9.0	3.1	3.0	0.95	0.13
Lognormal	46	Normal	101	9.0	2.9	3.0	0.96	0.11
Lognormal	117	Normal	68	9.0	2.9	3.0	0.96	0.11
Lognormal	117	Normal	101	9.2	3.1	3.0	0.94	0.14

[a] $\times 10^3$; $\alpha\beta \times 10^3 = 9.1$.

This was selected to be slightly larger than but consistent with the health effect estimates obtained in time series studies. Note that the health effect estimates from the simulations are all approximately half the true value for β. The attenuation parameter α is not identifiable from only the time series model. Since no accommodation for α has been made, the model estimates the quantity $\beta^e = \alpha\beta = 0.0091$. This is reflected in the very low coverage probabilities for $\beta = 0.0182$. It is also noteworthy that the variance estimates are constant across the range of distributional assumptions for parameters making up X_{it}^{N} and that the two approaches to estimating the variability in $\alpha\beta$ give approximately the same results.

These first simulations show that when ambient and nonambient source exposures are independent, changes in the distribution of the nonambient exposure have no impact on the health effect parameter estimate for the ecologic time series design. This shows that the well-known Berkson measurement error result holds, even though the nonlinear disease model allows for the possibility that specification bias may affect the estimates. In this particular example specification bias is absent because all the unmeasured variability in X_{it}^{P} is independent of the measured exposure X_t^{C}. Thus there is no correlation between $\text{var}(X_{it}^{\text{P}}|t)$ and X_t^{C}. Additional simulations with nonzero correlation (not shown) did not show notable evidence of specification bias, likely because of the near linearity of the health effect in this situation.

Role of the Attenuation Parameter α_{it}

The second set of simulations allows variation in the fraction of ambient concentration an individual is exposed to, α_{it}. Table 2 gives the results. The first two rows show the effect of α_{it} varying across the population with a constant mean of 0.5 on each day. The estimate of β^e and its standard error are unchanged with the added variation in α_{it} for both the uniform and normal distributions. For the remaining results we allow $\alpha_{it} \sim N(\alpha_t, 0.15)$ and then specify a function for α_t. (Note that the concentration series X_t^{C} is characterized by higher values in the winter and lower in the summer.) In the next set of three results, α_t varies over time as defined by $k\cos(180 + t/(365/\pi/2) - 0.056) + 0.5$ where t indexes day of year and $k = 0.1$, 0.2, 0.45, respectively. This constrains α_t to be low in the winter with progressively larger seasonal variation. The result is an estimate of β^e that is biased low with greater bias as the variation in α_t over time increases. Similarly for the last set of three results α_t varies over time as $k\cos(t/(365/\pi/2) + 0.056) + 0.5$. Here, α_t is high in the winter with progressively larger seasonal variation. The resulting estimates of β^e are biased high with increasingly more bias when α_t has the strongest seasonal structure. Our results show that variation in α_{it} matters in time series studies only to the

Table 2. Estimates of $\beta^e = \alpha\beta$ under variable α_{it}.

α_{it}, correlation with PM	$\widehat{\alpha\beta}$[a]	$\sqrt{\widehat{\text{var}(\widehat{\alpha\beta})}}$[a]	95% Cov.
U (0.15, 0.85), none	9.4	2.9	0.95
N (0.5, 0.15), none	9.2	3.0	0.96
Low winter, weak	8.4	3.0	0.96
Low winter, mod	7.4	3.1	0.92
Low winter, strong	5.4	3.1	0.80
High winter, weak	10.3	3.0	0.95
High winter, mod	11.1	3.0	0.89
High winter, strong	13.5	2.9	0.65

[a] $\times 10^3$; $\alpha\beta \times 10^3 = 9.1$.

degree that it is temporally correlated with ambient concentration.

Impact of Exposure Measurements

The third set of simulations evaluates which properties of an ambient measurement are important for ecologic time series studies. Since there will be some systematic differences within a geographic area of ambient concentrations, the common practice of using a single ambient monitor from a fixed location to represent population exposure needs study. As discussed in incorporating a measurement error model; the addition of a constant to the ambient concentration for each individual will not have much impact on the health effect estimate. With respect to random variation in the monitor measurements, in the first part of Table 3 we assume that there is some spatial variation in the ambient monitors and assess the effect of using only one monitor or the average of several monitors. The spatial variation is randomly distributed and relatively small (see previous section), adding error with variance 4 to the population ambient concentration, X_t^C. The strength of association between $\widehat{\overline{X_t^C}}$ and X_t^C is described by $\mathrm{corr}(\widehat{\overline{X_t^C}}, X_t^C)$. Using a single monitor ($m=1$) in the health effects analysis results in a small but noticeable attenuation. This disappears as the number of monitors over space increases.

In the second part of Table 3, we continue to assume that total personal exposure is the relevant exposure, and therefore use it directly in the time series regression model. The nonambient source variation is as given in (3) with both b_i and ε_{it} normally distributed with variances 34, 100, respectively. We use one or the average of up to $m=100$ individual exposures to represent the population exposure. (Note that $\overline{X_t^P} = \beta_0 + \alpha X_t^C$ so in the absence of measurement error attenuation β can be recovered directly from analysis of the time series model with $\overline{X_t^P}$ as the exposure. For comparability with estimates in the previous section of the table, we actually use $(\widehat{\overline{X_t^P}} - \beta_0)/\alpha$, in the analysis.) We describe the relationship between the personal exposure measurement and X_t^C as $\mathrm{corr}((\widehat{\overline{X_t^P}} - \beta_0)/\alpha, X_t^C)$. For comparison we also include the daily average of the personal exposures on all 100,000 members of the population. Note that there is a great deal of attenuation in the health effect estimate for small samples of personal exposure measurements on each day. To obtain a reasonably unbiased estimate of the exposure effect parameter, one would need a minimum of 100 personal exposures each day in this design, even though nonambient source exposure also drives the health effect. This study points out that for an aggregated exposure, exposure variation that is not associated with the aggregated outcome behaves like classical measurement error in the analysis. This is the case even though error in an aggregated exposure would typically be called Berkson error. In group-level studies, components of exposures that vary on the same scale as the outcome may be the most suitable exposure measurements for producing unbiased exposure effect estimates. Averaging can remove the noise associated with other sources of variation, but when these sources are quite variable, it is important to average over a large number of measurements in order to avoid attenuation bias.

Table 3. Estimates of $\beta^e = \alpha\beta$ when exposure is estimated from multiple monitors or people.

M	Corr w/X_t^C	$\widehat{\alpha\beta}$[a]	$\sqrt{\widehat{\mathrm{var}(\widehat{\alpha\beta})}}$[a]	95% Cov.
$\widehat{\overline{X_t^C}}$				
1	0.95	8.4	2.9	0.94
3	0.98	9.0	3.0	0.93
10	0.99	9.2	3.0	0.93
$(\widehat{\overline{X_t^P}} - \beta_0)/\alpha$				
1	0.29	0.9	0.9	0.00
3	0.47	2.2	1.5	0.01
10	0.69	4.8	2.2	0.48
50	0.91	7.7	2.5	0.95
100	0.95	8.5	2.7	0.96
100,000	0.99	9.3	2.8	0.96

[a] $\times 10^3$; $\alpha\beta \times 10^3 = 9.1$.

Implications of the Measurement Error Model in Time Series Studies

The simulation studies in this section indicate that the size of the population- and time-averaged attenuation parameter α affects the health effect estimates obtained and that by using X_t^C, time series studies estimate $\alpha\beta$. However, significant variation in nonambient exposure, or in ambient source exposure that is *independent* of ambient concentration does not bias the effect estimate $\widehat{\alpha\beta}$ when X_t^C is the exposure measurement. Thus, random variation in α_{it} across the population has no effect on the estimation of $\alpha\beta$. In contrast, shared population seasonal variation in α_{it} has a big impact on estimation of $\alpha\beta$. Since it is likely that α_{it} varies over time and across populations, it will be essential to better characterize this parameter. For instance, Janssen et al. (2002) discuss how city-specific estimates of $\alpha\beta$ vary by prevalence of air conditioning use, an important determinant of population-average ambient attenuation.

In contrast to time series studies that use ambient monitor measurements, it is not practical to use average personal exposure in time series studies when there is no shared variation in nonambient source exposure across the population over time. While average personal exposure allows the parameter directly estimated in the time series model to be β, unless there are measurements on hundreds of individuals on each day, the large variation in nonambient source exposure and its independence of ambient exposure effectively induces a huge amount of classical measurement error when average personal exposure is used as the predictor in the time series

design. Zeger et al. (2000) suggest that risk-weighted average personal exposure is the desired exposure measurement in time series studies. Our work supports that conclusion from the perspective of directly estimating the target parameter β, but differs by further showing that the nonambient source exposure variability will induce a large amount of classical measurement error into the analysis (unless the daily samples for estimating $\overline{X_t^P}$ are huge). Using the average of multiple ambient monitor measurements in the time series regression gives estimates of $\alpha\beta$ that are less biased from classical measurement error than risk-weighted measurements of $\overline{X_t^P}$ give for β.

Discussion

Understanding the interplay between exposure, the parameters of interest, and the study design is a necessary part of study planning to achieve the goal of estimating the health effects of environmental exposures. Meaningful progress in estimating environmental exposure effects is possible when the properties of the exposure effect parameters are understood and estimable. Key study design components include: identifying the exposure metric and its main sources of variation, the target population, the level of analysis, and the sampling plan.

A panel study design with exposure variation both within and between individuals allows one to estimate separate effects of exposure that can have different interpretations. While the assumption of no contextual effects means the underlying parameters are identical, their estimates can be different due to influences of exposure measurement error and confounding. Sources of exposure variation that are collapsed (or nearly collapsed) in an aggregated analysis will provide little, if any, information for estimation.

The impact of study design or the loss of information due to aggregation depends upon how exposure varies for a given disease model. Understanding exposure variation in the context of the specific study is necessary even when the exposure is treated as fixed in the statistical analysis. For air pollution exposure, the assumed independence of ambient and nonambient exposure sources and the ready availability of ambient concentration measurements have important implications for selection of designs that are most effective for estimating health effects.

Our simulation studies of the effect of exposure and exposure measurements in the time series design are based upon assumptions supported by the published literature and/or our analyses of a large multi-year panel study in Seattle. While we believe these assumptions are reasonable, some are not fully evaluated or replicated. One of our key assumptions is that the ambient and nonambient source exposures are independent. When this assumption holds, the time series design only uses information about the ambient source exposure. We found that time series studies recover a function of the relative risk parameter fairly well in practice because the ambient concentration measurements do a reasonable job of summarizing time-varying population average concentration, which forms the basis of ambient source exposure, and the design can take advantage of an entire population of a geographic area.

For time series studies, the sources of measurement error can be quantified as follows. The first is random variation in the concentration measurement from the ambient concentration for the population on a day. Our simulations showed that (at least for $PM_{2.5}$) much of the effect on the health effect estimates due to the first source of error can be corrected by using the daily average of a few monitors. A second important source of error has to do with the temporal variation in the population average fraction of ambient exposure (α_t). This can result in strong positive or negative biases, depending on how α_t is correlated with seasonal trends in ambient concentration. As a third source of error, the difference between average personal exposure and ambient measurements is important only to the degree that one is interested in estimating the relative risk parameter β directly from the time series studies. Time series studies estimate $\alpha\beta$. Since $\Delta X_t^C \alpha\beta = \Delta X_t^A \beta$ (where Δ means change), the same relative risk parameter β is incorporated regardless of whether the model is parameterized in terms of ambient exposure or ambient concentration. When $\overline{X_t^P}$ is used as the exposure measurement, β can be estimated directly. However, when only a few individuals provide data to estimate $\overline{X_t^P}$, the model will suffer from severe classical measurement error bias. Under the assumption of independence of ambient and nonambient source exposures, time series studies target ambient source exposure effects and thus do not need to incorporate average personal exposure measurements. Interpretation of these studies would benefit from better understanding of whether and how the attenuation of ambient concentration (α_t) covaries with concentration (X_t^C) and whether this covariance varies across areas.

Air pollution time series studies have been very successful because they use routinely collected data, including ambient air pollution monitoring and health surveillance measures. These data are publicly available and thus inexpensive and fairly easy to obtain. The nature of variation of pollution exposure, particularly $PM_{2.5}$ (which has relatively little variation over space), means that the time series design can estimate the health effect due to $PM_{2.5}$ exposure even though ambient source exposure is a small component of the variation due to total personal exposure.

We suggest multiple research paths to better understand the role of exposure, measurements, and study design on estimating health effects. Attention should be broadened to other study designs, particularly designs that target the the health effects of chronic exposures. Within any single study design, much more work is needed to assess the validity of the

basic exposure model and understand how exposure varies with respect to it. For $PM_{2.5}$ we need to continue to assess the validity of the assumption that ambient and nonambient source exposures are independent. We need to characterize the distribution of α_{it} and determine how it varies over time. Further work to assess the spatial and spatiotemporal distribution of ambient exposure is also necessary. This research should be repeated for other PM size fractions and gaseous pollutants. Comparisons across pollutants, for instance in the degreee of spatial heterogeneity in the ambient source or the degree of attenuation of ambient concentration for personal exposure, may give good insight into the relative performance of the time series estimates for these pollutants.

Acknowledgements

This work was supported by Grants ES08062-03 from the National Institute of Environmental Health Sciences, NIH and R827355 from the US Environmental Protection Agency. The contents are solely the responsibility of the authors and do not necessarily represent the official views of NIEHS or EPA.

References

Allen R., Larson T., Wallace L., Sheppard L., and Liu L.-J.S. Use of real-time light scattering data to estimate the contribution of infiltrated and indoor-generated particles to indoor air. *Environ Sci Technol,* 2003: 37: 3484–3492.

Burke J.M., Zufall J.J., and Ökaynak H. A population exposure model for particulate matter: case study results for PM2.5 in Philadelphia, PA. *J Expos Anal Environ Epidemiol* 2001: 11: 470–489.

Carroll R.J., Ruppert D., and Stefanski L.A. *Measurement Error in Nonlinear Models*. Chapman & Hall, New York, 1995.

Clayton D. Models for the longitudinal analysis of cohort and case–control studies with inaccurately measured exposures. In: Dwyer, J.H., Feinleib, M., Lippert, P., Hoffmeister, H. (Eds.). *Statistical Models for Longitudinal Studies of Health of Monographs in Epidemiology and Biostatistics Vol. 16*. Oxford University Press, Oxford, 1991, pp. 301–331, Chapter 12.

Dominici F., Daniels M., Zeger S.L., and Samet J.M. Air pollution and mortality: estimating regional and national dose-response relationships. *J Am Statist Assoc* 2002: 97(457): 100–111.

Dominici F., Zeger S.L., and Samet J.M. A measurement error model for time-series studies of air pollution and mortality. *Biostatistics* 2000: 12: 157–175.

Duan N. Stochastic microenvironmental models for air pollution exposure. *J Expos Anal Environ Epidemiol* 1991: 1(2): 235–257.

Firebaugh G. A rule for inferring individual-level relationships from aggregate data. *Am Sociol Rev* 1978: 43: 557–572.

Gilks W.R., and Richardson S. Analysis of disease risks using ancillary risk factors, with application to job-exposure matrices. *Statist Med* 1992: 11: 1443–1463.

Janssen N.A.H., Schwartz J., Zanobetti A., and Suh H.H. Air conditioning and source-specific particles as modifiers of the effect of PM10 on hospital admissions for heart and lung disease. *Environ Health Perspect* 2002: 110(1): 43–49.

Koutriakis P., Briggs S., and Leaderer B. Source apportionment of indoor aerosols in suffolk and onondaga counties, New York. *Environ Sci Technol* 1992: 26: 521–527.

Kunzli N., and Tager I.B. The semi-individual study in air pollution epidemiology: a valid design as compared to ecologic studies. *Environ Health Perspect* 1997: 105: 1078–1083.

Liu L.-J.S., Box M., Kalman D., Kaufman J., Koenig J., Larson T., Lumley T., Sheppard L., and Wallace L. Exposure assessment of particulate matter for susceptible populations in Seattle, WA. *Environ Health Perspect* 2002: 111: 909–918.

Mage D., Wilson W., Hasselblad V., and Grant L. Assessment of human exposure to ambient paticulate matter. *J Air Waste Manage Assoc* 1999: 49: 1280–1291.

National Research Council Committee on Research Priorities for Airborne Particulate Matter. Research Priorities for Airborne Particulate Matter: Immediate Priorties and a Long-Range Portfolio, National Academy Press, Washington, DC, 1998.

National Research Council Committee on Research Priorities for Airborne Particulate Matter. Research Priorities for Airborne Particulate Matter: III Early Research Progress, National Academy Press, Washington, DC, 2001.

Ott W., Wallace L., and Mage D. Predicting particulate (PM_{10}) personal exposure distributions using a random component superposition model. *J Air Waste Manage Assoc* 2000: 50: 1390–1406.

Piantadosi S., Byar D.P., and Green S.B. The ecological fallacy. *Am J Epidemiol* 1988: 127(5): 893–904.

Salway R., and Wakefield J. Sources of bias in ecological studies of non-rare events. *Environmental-Ecological Statistics*, in press.

Schwartz J., Laden F., and Zanobetti A. The concentration-response relation between PM2.5 and daily deaths. *Environ Health Perspect* 2002: 110: 1025–1029.

Sheppard L. Ecologic study design. In: El-Shaarawi, A.H., Pierorcsh, W.W. (Eds.). *Encyclopedia of Environmetrics, Vol. 2*. John Wiley and Sons, New York, 2002, pp. 602–606.

Sheppard L. Insights on bias and information in group-level studies. *Biostatistics* 2003: 4: 265–278.

Sheppard L., and Damian D. Estimating short-term PM effects accounting for surrogate exposure measurements from ambient monitors. *Environmetrics,* 2000: 11: 675–687.

Sheppard L., and Prentice R.L. On the reliability and precision of within and between population estimates of relative rate parameters. *Biometrics* 1995: 51: 853–863.

Wilson W.E., Mage D.T., and Grant L.D. Estimating separately personal exposure to ambient and nonambient particulate matter for epidemiology and risk assessment: why and how. *J Air Waste Manage Assoc* 2000: 50: 1167–1183.

Yu O., Sheppard L., Lumley T., Koenig J.Q., and Shapiro G. Effects of ambient carbon monoxide and atmospheric particles on asthma symptoms: results from the CAMP air pollution ancillary study. *Environ Health Perspect* 2000: 108: 1209–1214.

Zeger S.L., Thomas D., Dominici F., Samet J.M., Schwartz J., Dockery D., and Cohen A. Exposure measurement error in time-series studies of air pollution: concepts and consequences. *Environ Health Perspect* 2000: 108: 419–426.

Zidek J.V., Jeloche J., Shaddick G., Chatfield C., and White R. A computational model for estimating personal exposure to air pollutants with application to London's PM10 in 1997. Technical Report 2003-3, Statistical and Applied Mathematical Sciences Institute, RTP, NC,, 2003.

Zidek J.V., Meloche J., Le N.D., and Sun L. Combining statistical and Computer models for health risk assessment (exposure analysis). In: Nex-Antn, V., Feffeirra, E. (Eds.). *Statistical Modeling*. Proceeings of the 15th International Workshop on Statistical Modelling: New Trends in Statistical Modelling, Universidad del Pais Vasco, Bilbao, 2000, pp. 95–106.