

**Exposure Measurement Error in Time-Series Studies of Air Pollution:
Concepts and Consequences**

**S. L. Zeger¹, D. Thomas², F. Dominici¹, J. Samet¹, J. Schwartz³,
D. Dockery³, and A. Cohen⁴**

Johns Hopkins University¹
School of Public Health
615 North Wolfe Street
Baltimore, MD 21205

Department of Preventive Medicine²
University of Southern California School of Medicine
1540 Alcazar Street, Suite 220
Los Angeles, CA 90033

Harvard University³
665 Huntington Avenue
Boston, MA 02115

Health Effects Institute⁴
955 Massachusetts Avenue
Cambridge, Massachusetts 02139

Correspondence: Dr. Scott Zeger¹

Keywords: measurement error, air pollution, time series, exposure, design methods

Abstract

Misclassification of exposure is a well-recognized, inherent limitation of epidemiologic studies of disease and the environment. For many agents of interest, exposures take place over time and in multiple locations; accurately estimating the relevant exposures for an individual participant in epidemiologic studies is often daunting, particularly within the limits set by feasibility, participant burden, and cost. Researchers have taken steps to deal with the consequences of measurement error by limiting the degree of error through a study's design, estimating the degree of error using a nested validation study, and by adjusting for measurement error in statistical analyses.

In this paper, we address measurement error in observational studies of air pollution and health. Because measurement error may have substantial implications for interpreting epidemiologic studies on air pollution, particularly the time series analyses, this paper sets out a systematic conceptual formulation of the problem of measurement error in epidemiologic studies of air pollution and considers the consequences within this formulation. When possible, we have used available, relevant data to make simple estimates of measurement error effects.

We start with an overview of measurement errors in linear regression, distinguishing two extremes of a continuum-Berkson from classical type errors, and the univariate from the multivariate predictor case. We then propose one conceptual framework for the evaluation of measurement errors in the log-linear regression used for time series studies of particulate air pollution and mortality, identifying three main components of error. We present new, simple analyses of data on exposures of

particulate matter less than 10 μm in aerodynamic diameter (PM_{10}) from the PTEAM (Particle Total Exposure Assessment Methodology) study. Finally, we summarize open questions regarding measurement error and suggest the kind of additional data necessary to address them.

1. Introduction

Misclassification of exposure has long been recognized as an inherent limitation of epidemiologic studies of the environment and disease (1). For many agents of interest, exposures take place over time and in multiple locations so that it is difficult to accurately estimate the relevant exposures for individual study participants, particularly within the limits set by feasibility, participant burden, and cost. In general, exposure measurement error tends to blunt the sensitivity of epidemiologic studies for detecting effects of environmental agents, although the specific impact of exposure error on effect estimates depends on several factors including the study design, the types of error, and the relationships between the outcome and the independent variables (1, 2). As the problem of exposure error has become well-recognized, researchers have taken steps to control its consequences by limiting the degree of error through careful study design and data collection, by estimating the degree of error using a nested validation study, and by making adjustments for measurement error in statistical analyses.

In this paper, we address the problem of exposure error in observational, ecologic time-series studies of air pollution and health. Pollution of outdoor air is a public health concern throughout the world. For decades, epidemiologic studies have been a cornerstone of our approach to investigating the health effects of air pollution and have been a principal basis for setting regulations to protect the public against adverse health effects. Two broad types of observational study designs have been used in research on air pollution: ecological or aggregate-level studies, either cross-sectional or time-series in design, and individual-level studies, primarily of the cross-sectional or cohort designs. In ecological studies, population-level indicators of

exposure are typically drawn from centrally-sited air pollution monitors. In individual-level cross-sectional and cohort studies, exposure estimates for individual participants may be based on centrally located monitors, on the combination of central monitors with personal records of environments where participants spend time, or on personal exposure monitoring (3).

Regardless of study design, any pollution exposure assessment strategy will introduce some degree of exposure measurement error. For example, in the Six Cities Study, a prospective cohort study of air pollution and respiratory health and mortality, exposure estimates for persons from each of the six cities were based on centrally sited monitors (4, 5). Exposures were further characterized for samples of participants using personal monitors and monitors placed in their homes; the resulting data provide an understanding of the components of error associated with using the central site data for all participants.

The problem of measurement errors in predictor variables in regression analysis has been carefully studied in the statistics and epidemiological literature for several decades. Fuller (6) summarizes early research on linear regression with so-called “errors-in-x” variables. Carroll et al. (7) extend this literature to generalized linear models including Poisson, logistic, and survival regression analyses. Thomas and colleagues (2) presents an overview of the exposure error or “misclassification problem” from the general epidemiological perspective. For recent illustrations of statistical approaches to measurement error in epidemiological research, see Spiegelman et al. (8), Willett (9), and Pierce et al. (10).

In one of the early papers on the topic of exposure error in studies of air pollution, Shy and colleagues (11) described the problem and addressed its consequences in an epidemiologic framework. Goldstein (12, 13) recognized that a single monitoring station may not adequately represent a geographic area and conducted an analysis of correlations among concentration data from several monitors in New York City. In the ensuing decades, there has been deepening understanding of measurement error in general and its potential implications for the study of air pollution (14, 15).

During the 1990s, substantial new evidence, largely from ecologic, time-series analyses of air pollution and mortality, showed that daily variation in ambient measures of particulate air pollution, within current standards of the US Environmental Protection Agency, was associated with daily mortality levels (16). Strong concerns have been raised about interpreting these associations in view of potential errors in the exposure measurements. In a series of papers, Lipfert and Wyzga (17-19) have suggested that the central monitoring data used in the time-series analyses have an uncertain relationship with exposures of individuals in the study communities; they have further argued that those errors vary among pollutants, complicating interpretation of any multi-pollutant models. Lipfert and Wyzga have referred specifically to an analysis by Schwartz, Dockery, and Neas (20) which attributed effects on mortality to fine rather than coarse particles, based in part on the results of multivariable models which included variables for both particulate measures.

A number of exposure assessment studies have found sizable differences between actual personal exposures to particles and estimates based on central monitor values (21). Some have questioned whether the observed associations are plausible

given these findings. However, Schwartz, Dockery, and Neas have responded that as the number of deaths per day is calculated over the population, the relevant exposure measure is the *mean* of personal exposures on that day, which is probably more tightly correlated with central station monitoring than individual exposures. Janssen and colleagues (22) have reported that much of the variation in PM₁₀ measurements is between people and that the longitudinal correlation between average and ambient PM₁₀ measures is relatively high. The debate over measurement error and its consequences has taken place, however, without the development of a more comprehensive formulation of the problem.

Because exposure measurement error may have substantial implications for interpreting epidemiologic studies on air pollution, particularly the time-series analyses, this paper describes one systematic conceptual formulation of the problem of exposure error in epidemiologic, time-series studies of air pollution and considers the possible consequences for relative risk estimation. We have used available and relevant data to obtain rough estimates of the magnitudes of the effects of measurement error for one city.

Section 2 presents an overview of the main ideas on exposure measurement errors in linear regression, distinguishing Berkson from classical type errors and the univariate from multivariate predictor cases. Section 3 develops a conceptual framework for evaluation of measurement errors in the log-linear regression models used for time-series studies of particulate air pollution and mortality, identifying three main components of error. In Section 4.1, we present new analyses of data on exposures to particulate matter less than 10 μm in aerodynamic diameter (PM₁₀) from

the PTEAM (Particle Total Exposure Assessment Methodology) study (23). We illustrate how data on personal and ambient exposure levels can be used to assess the effects of measurement error on the estimated associations of PM₁₀ with daily mortality. In Section 4.2, we illustrate a statistical approach for assessing the bias in a relative risk estimate caused by exposure measurement error. Section 5 summarizes the open questions regarding measurement error and proposes the additional data needed to address more effectively these questions.

2. Overview of Measurement Error Effects in Regression Models

This section sets out the fundamental concepts of how exposure error can affect an epidemiologic study of pollution and health. We do so by considering the effects of exposure measurement error in a standard linear gaussian regression model. This topic has been treated in full detail elsewhere (2, 6, 7, 24, 25). For simplicity, consider a regression of the health response y_t (e.g., log mortality rate on day t) and predictors x_t (e.g., PM₁₀, O₃, weather):

$$y_t = \mathbf{a} + \mathbf{b}_x x_t + \mathbf{e}_t \quad (\text{E1})$$

where \mathbf{a} and \mathbf{b}_x are regression coefficients to be estimated, and \mathbf{e}_t represents residual error which is assumed to be independent of x_t . Here \mathbf{b}_x is the expected change in mortality per unit change in true exposure. Given observations (x_t, y_t) , $t = 1, \dots, T$ and appropriate assumptions about the distribution of the residuals, ordinary least-squares estimation provides optimal (unbiased and minimum variances) estimates of the regression coefficients.

Now we assume that instead of the true exposure levels x_t , we have only an imperfect measure of exposure, denoted z_t . The overall difference between x_t and z_t comprises multiple components of error including differences: between individual- and population-average exposures; between population-average exposures and ambient levels at central sites; and between actual ambient levels and the measurements of those levels. Suppose we regress the health outcome y_t on the imperfect z_t rather than x_t which is unavailable:

$$y_t = \mathbf{a}^* + \mathbf{b}_z z_t + \mathbf{e}_t^*$$

How will $\hat{\mathbf{b}}_z$ differ from $\hat{\mathbf{b}}_x$?

To answer this question, we will first assume that z_t is a “surrogate” for x_t which means that, given x_t , there is no additional information in z_t about y_t . We then can distinguish two fundamentally distinct types of relationships between the true and measured exposures, which represent poles of a measurement error continuum. The first type is referred to as the “classical error model” (7) in which we assume that z is an imperfect measure of x , so that the average z within each x stratum equals x ($E(z|x) = x$). Then it follows that the measurement error $z - x$ is uncorrelated with the true value x . This “classical model” is a reasonable one for the difference between measured ambient levels of pollution and the true values for a measuring device that is unbiased. That is, when the true level of pollution is x , an unbiased instrument will measure x on average, even if individual measurements z differ from x .

The second type of model for measurement error is the “Berkson error model” (2). In this model, we assume that the average value of the true exposure x within

each stratum of measured level z equals z ($E(x|z) = z$). This Berkson model is appropriate when z represents a measurable environmental factor that is shared by a group of participants whose individual exposures x might vary because of time-activity patterns. For example, z might be the spatially averaged ambient level of a pollutant without major indoor sources and x might be the personal exposures, which when averaged across people, match the ambient level.

“Classical” and “Berkson” models for exposure measurement errors represent two extremes of a continuum. Most exposure errors combine elements of each, but because the consequences on risk assessment of classical and Berkson errors differ, it is useful to consider each in turn. In the Berkson error case, if we regress y_t on z_t , rather than on x_t , the estimate \hat{b}_z is an unbiased estimate of the coefficient b_x , which would be obtained by regressing y_t on the actual exposure x_t . That is, having z_t rather than x_t does not lead to bias in the regression coefficients under the surrogacy assumption. The exposure measurement error does increase the variance of the regression coefficient, however, since having z_t rather than x_t is obviously not as informative about the coefficient b_x . Bias is not introduced, however. The same is true if the average x at each value of z differs from z by a fixed amount a , i.e.,

$$E(x|z) = z - a.$$

In contrast, under the classical error model, \hat{b}_z obtained by regressing y_t on the imperfect measure exposure z_t , is a biased estimate of b_x . In the simple linear regression with one explanatory variable, \hat{b}_z is expected to be smaller than b_x or “attenuated”. The degree of attenuation increases as the variance of the exposure error

increases. Again, a constant difference in the expected values of the two measures does not change this result.

It is straightforward to establish the results summarized above on the effects of exposure error on simple linear regression coefficients and useful to do so in advance of considering a multiple regression case. To re-establish notation, the model of interest is

$$y_t = \mathbf{a} + \mathbf{b}_x x_t + \mathbf{e}_t \quad (\text{E1})$$

but because x_t is unobserved we instead might regress y_t on z_t :

$$y_t = \mathbf{a}^* + \mathbf{b}_z z_t + \mathbf{e}_t^* \quad (\text{E2})$$

The question is how will $\hat{\mathbf{b}}_z$ from (2) estimate \mathbf{b}_x in (E1). Under the classical error,

$E(z_t | x_t) = x_t$ z_t is assumed to vary about x_t , that is so that by (2),

$$E(y_t | x_t) = \mathbf{a}^* + \mathbf{b}_z E(z_t | x_t) = \mathbf{a}^* + \mathbf{b}_z x_t \quad (\text{E3})$$

Comparing (E3) and (E1) shows that $\mathbf{b}_z = \mathbf{b}_x$ in the Berkson error case; that is $\hat{\mathbf{b}}_z$ is an unbiased estimate of \mathbf{b}_x . Adding a constant to one exposure variable only affects the intercept.

Under the classical model, x_t is assumed to vary about z_t or $E(x_t | z_t) = z_t$ which does not imply $E(z_t | x_t) = x_t$. If we further assume that x_t and $z_t - x_t$ are jointly normally distributed, it can be shown that

$$E(y_t | x_t) = \mathbf{a}^{**} + c \mathbf{b}_x x_t$$

where c is an attenuation factor between 0 and 1 given by $c = \text{var}(x_t) / (\text{var}(x_t) + \text{var}(\mathbf{d}_t))$ where $\mathbf{d}_t = z_t - x_t$ is the exposure error. Again, a constant difference between the two exposure measures only changes the intercept.

Thus, the estimated regression coefficient is biased towards zero. In one pertinent case, $\mathbf{b}_x = 0$, the naive estimate $\hat{\mathbf{b}}_z$ is unbiased with $E(\hat{\mathbf{b}}_z) = \mathbf{b}_x = 0$; that is, under the classical error model, measurement error does not lead to spurious associations if there is truly no association. Random variation, of course, can produce such associations by chance, as it can absent measurement error. However, the probability of such false positive associations (the Type I error rate) remains the same.

For estimating effects of air pollution on mortality, realistic models have elements of both classical and Berkson error models. In general, the effect of such exposure errors is intermediate between the two extreme models. The effect of measurement error, therefore likely depends upon the direction and magnitude of the correlation of measurement errors with the measured exposures and not just upon the variance of the measurement errors.

We now turn to the more complex situation of multi-pollutant models. Such models are often applied in an attempt to estimate the independent effect of a pollutant present in a mixture with other pollutants. For example, in an analysis of air pollution and mortality in Philadelphia, Kelsall et al. (26) regress mortality on as many as five pollutants. Because little empirical evidence about the simultaneous errors in multiple pollutants is currently available, this section only lays a foundation that can inform the design of future studies, as discussed in Section 5. Confining attention to the classical and the Berkson error cases, we again assume a linear regression model of the form given by Equation 1, where x_t now represents a vector of exposure variables, with a corresponding vector of regression coefficients \mathbf{b}_x , and z_t denotes a vector of measurements of each exposure variable. In the Berkson error case, the assumption

that x_t is an imprecise version of z_t or $E(x_t|z_t) = z_t$ still assures that the estimates of the regression coefficients are unbiased, as in the univariate instance. But under the classical error model, the multiple regression extension is not so straightforward. As before, we assume that z_t is an imprecise measure of x_t , i.e., $E(z_t|x_t) = x_t$. To compute $E(x_t|z_t)$ the average x_t at each z_t , let V denote the covariance matrix of x_t and let T denote the covariance matrix of the difference $\mathbf{d}_t = z_t - x_t$, V be the variance of x_t and, as before, we assume that δ and x are independent. Then, the matrix generalization of the earlier result is that $\mathbf{b}_z = \mathbf{b}_x C$ where $C = T(T+V)^{-1}$. Now it is no longer true that $\mathbf{b}_{z_j} < \mathbf{b}_{x_j}$ for each component (j) and estimates of regression coefficients can be biased toward or away from the null; that is, positive associations can be produced even though the true coefficient for a particular component is zero, when the component is correlated with at least one component having a non-zero effect.

Table 1 illustrates the magnitude of bias that can result from regressing y_t on two predictors z_{1t} and z_{2t} instead of on x_{1t} and x_{2t} . This example might refer to estimating the effects of PM₁₀ and O₃ on mortality when ambient values (z s) instead of personal exposure (x s) are available. We assume $z_{1t} = x_{1t} + \mathbf{d}_{1t}$ and $z_{2t} = x_{2t} + \mathbf{d}_{2t}$, $V_{11} = \text{Var}(x_{1t}) = V_{22} = \text{Var}(x_{2t}) = 1$. The table presents the expected values for the estimated regression coefficients when the true values are both one ($\mathbf{b}_{x_1} = \mathbf{b}_{x_2} = 1$) for varying values of the correlation between x_{1t} and x_{2t} , the variances of \mathbf{d}_{1t} and \mathbf{d}_{2t} , and the correlation between the measurement errors \mathbf{d}_{1t} and \mathbf{d}_{2t} . At present, there is little empirical evidence about the nature or size of the correlations between pairs of pollutant

measurements and the table is intended to illustrate the consequences of measurement error in the two-predictor model.

The first line of the table refers to an example in which there is no correlation between x_{1t} and x_{2t} and there is equal variability of the two exposure errors d_1 and d_2 , and these errors are not correlated; that is, the error in one predictor does not predict the error in the other. Here, there is an equal degree of attenuation in the coefficients for the two variables. With unequal variances, but no correlation, i.e., the sixth row, the degree of attenuation is greater for the variable with greater variance. If the exposures are correlated, but the errors are uncorrelated (the second and third rows), the two effect estimates are similarly altered with the direction of the effect depending on the sign of the correlation. Introducing correlation between the errors, i.e., the fourth and fifth rows, has an effect that depends on the pattern of correlation. The bottom half of **Table 1** shows more complex patterns with differing patterns of correlation and variation of the two errors. Some of the scenarios introduce substantially different effects of the two variables, but none yield effect estimates above the true value of one, even with more extreme differences in error variances or the two correlations.

Table 2 also addresses the consequences of measurement error in a two-variable model, but in this example only one variable (x_2) has a true effect; the other exposure x_1 has no effect on the health outcome y . Either correlation between x_{1t} and x_{2t} or their errors can introduce an apparent effect of x_1 on y . Some scenarios of variance and correlation even bring the apparent effects of the two variables quite close (e.g., the tenth and eleventh rows), but in every case, including more extreme situations

than shown, the estimate for the true predictor (b_2) is always larger than for the null predictor (b_1).

Some general conclusions can be offered concerning multi-pollutant models under this simple, classical error model.

- C1. There is a general tendency for the coefficient from the regression on z_t to be attenuated (smaller than) the corresponding coefficient from the regression on x_t , i.e., $b_{zj} < b_{xj}$ if all $b_{xj} > 0$.
- C2. The degree of attenuation of each coefficient depends, in large part, on its measurement error variance relative to the variance of the true exposure – i.e., T_{jj}/V_{jj} . Thus, the coefficients for variables that are measured with considerable error will be attenuated more than those of variables with less error.
- C3. Depending on the correlation structure of the attenuation matrix C , some of the effect of one variable, b_{xj} , may be transferred to the estimate of another variable's effect, \hat{b}_{zk} . Such transfers of effect are generally from a more poorly measured variable to a better measured variable. However, for such transfers to be large, the true exposure variables or their measurement errors need to be substantially correlated.
- C4. As a consequence of conclusion (C3), the estimate of a parameter can be biased away from the true value. However, this type of bias generally arises only with a very strong negative correlation between the measurement errors (e.g. the ninth-eleventh rows of **Table 2**).

- C5. Also as a consequence of (3), there will generally be spurious associations for a variable x_j which, in fact, has no effect only if x_j is substantially correlated with one or more variables which actually have an effect. Generally, the correlation amongst the errors has a larger influence on the bias than the correlation amongst the true pollutant levels.

These conclusions are obtained from and therefore pertain to the classical linear regression model with two predictors, assuming z_t is a surrogate for x_t (non-differential errors). The actual exposure measurement situation in the air pollution-mortality context is obviously more complex. First, log-linear, not linear, models are used, although the degree of non-linearity is usually small in mortality studies. Second, the measurement errors are not purely of the classical, non-differential type. For example, the degree of error for gaseous pollutants may depend on temperature or other covariates. Finally, errors may be multiplicative rather than additive. Nonetheless, the linear regression with classical measurement error is a leading case that provides insight into the major possible consequences of exposure errors.

3. Framework for Assessing Measurement Error Effects in Pollution-Mortality Studies

The previous section has set out fundamental concepts underlying statistical models of exposure measurement error. This section builds upon these concepts but focuses on the specific log-linear regressions used for assessing the pollutant-mortality association, controlling for weather variables. We identify three major components of measurement error and present a statistical framework for evaluating their potential

effects on the estimated pollutant-mortality associations. The discussion below is based upon the premise that the ideal investigation of the health effects of air pollution would be conducted at the individual level with measurements of personal exposure to pollutants. However, exposure and mortality data are generally only available after aggregation to a municipal level; little or no data from indoor air monitoring are available. Finally, air pollutant measurements are imprecise and this imprecision has consequences for estimates of pollutant effects on mortality, as described in Section 2 above.

To investigate the effects of exposure error in the log-linear regressions widely used to assess the pollutant mortality association, consider the following model for an individual's risk of mortality:

$$I_{it} = I_{0it} \exp(x_{it} \mathbf{b}_x) \quad (\text{E4})$$

where I_{it} is the risk of death for person i on day t , I_{0it} is that individual's baseline risk in the absence of exposure, ie $x_{it} = 0$, and $\exp(x_{it} \mathbf{b}_x)$ is the relative risk of death associated with the explanatory variables x_{it} . Let $y_{it} = 1$ if person i dies on day t and 0 if he does not. We typically observe the total number of deaths for a population $y_t = \sum_{i=1}^{n_t} y_{it}$, where $n_t \approx n$ is the population size on day t . By (E4), the expected total numbers of deaths I_t in a community is

$$I_t = E y_t = \sum_i I_{it} = \sum_i I_{0it} \exp(x_{it} \mathbf{b}_x). \quad (\text{E5})$$

In analyzing population-level data on mortality and air pollution, log-linear regressions of the following form have been fit

$$I_t = \exp(s(t) + z_t \mathbf{b}_z + u_t \mathbf{b}_u) \quad (\text{E6})$$

where $s(t)$ is an arbitrary but smooth function of time introduced to control for the confounding of longer-term trends and seasonality, z_t is the average of multiple monitor measurements of ambient pollution measurement for day t , and u_t are other possible confounders such as temperature and dew point temperature on the same and previous days.

If the regression coefficient \mathbf{b}_x for a pollutant in the personal risk model (E4) is the target for inference, how closely do estimates of \mathbf{b}_z from model (E6) approximate \mathbf{b}_x ? Below, we identify potential sources of bias in $\hat{\mathbf{b}}_z$ as an estimate of \mathbf{b}_x , using the concepts of Berkson and classical measurement error summarized in Section 2.

Figure 1 poses a model of the relationship between the personal exposure to a pollutant x_{it} for person i on day t and the available ambient values z_t measured with error by monitors. Assuming, for simplicity, a high degree of spatial homogeneity in ambient levels, personal exposure is contributed to by z_t^* , the true outdoor level and w_{it} , the indoor level which is also influenced by z_t^* from penetration of the pollutant in outdoor air into indoor spaces. For example, personal exposure to PM_{10} is determined by the time spent outdoors, the concentration during that time, and by the concentrations in indoor environments that are determined by indoor sources such as cigarette smoking and the penetration of particles indoors, as air is exchanged between the outdoors and the indoor environments. **Figure 1** further shows that personal risk of dying is influenced by a person's baseline risk in addition to the unobserved personal

exposure to pollutant x_{it} . Only the measured ambient pollution data are observed and are therefore shown in a rectangular box.

In considering the consequences for \hat{b}_z as an estimate of b_x of having an imprecise measure of ambient pollution z_t , rather than actual personal exposure x_{it} , it is useful to begin by decomposing the pollution measurement difference between x_{it} and z_t into three components:

$$x_{it} = z_t + (x_{it} - \bar{x}_t) + (\bar{x}_t - z_t^*) + (z_t^* - z_t). \quad (\text{E7})$$

Here, $(x_{it} - \bar{x}_t)$ is the error due to having aggregated rather than individual exposure data; $(\bar{x}_t - z_t^*)$ is the difference between the average personal exposure and the true ambient pollutant level; and $(z_t^* - z_t)$ represents the difference between the true and the measured ambient concentration.

The first term $(x_{it} - \bar{x}_t)$ is an example of Berksonian error so that in a simple linear model, having aggregate rather than individual exposure does not itself lead to bias into the regression coefficient. The second term $(\bar{x}_t - z_t^*)$ is not Berksonian and is likely to be a source of bias. The final term $(z_t^* - z_t)$ is largely of the Berkson type if the average of the available monitors z_t is an unbiased estimate of the true spatially

*
averaged ambient level z_t^* .

We can now further study the effects of these three terms on risk estimation by substituting the decomposition in (E7) into Equation E5. After some straightforward calculations detailed in the Appendix, the expected number of deaths on day t can be written

$$Ey_t = \exp\left(\log(n_t \bar{I}_{0t}) + z_t \mathbf{b}_x + [(\bar{x}_t^{(w)} - \bar{x}_t) + (\bar{x}_t - z_t^*) + (z_t^* - z_t)] \mathbf{b}_x\right). \quad (\text{E8})$$

Here \mathbf{b}_x is the personal log-relative risk of interest from Equation E5. Note the approximation (E8) retains only linear terms in the expansion of an exponential function. The second order terms, are an order of magnitude smaller and are ignored to simplify the exposition. For studies of particulate pollution effects on mortality, the effect sizes are on the order of a percent or two so that ignoring second order terms should not qualitatively affect the results. In studies of morbidity, higher order terms may be more important.

The total baseline risk ($n_t \bar{I}_{0t}$) almost certainly varies smoothly over time, since it is an average risk over a large population. Hence, it will be appropriately controlled for in log-linear regressions by inclusion of the smooth $s(t)$ in Equation (E6). We now consider $z_t \mathbf{b}_x$ and the three components of error in turn.

The first error term $\bar{x}_t^{(w)} - \bar{x}_t$ is proportional to the difference between the baseline-risk weighted average personal exposure and the unweighted average personal exposure. It derives from the Berkson error ($x_{it} - \bar{x}_t$) and produces no bias in the linear, unaggregated model. This difference due to risk weighting in our log-linear model with person-specific baseline risks is likely to be small and to vary slowly over time. Hence, it can be adequately controlled by inclusion of the smooth function $s(t)$ in the log-linear regression of y_t on z_t . One scenario in which this difference would vary from day to day and therefore not be adequately controlled would occur if the more frail individuals were to follow pollution reports (or a correlate like weather) and reduce their exposures to ambient air on high pollution days by, for example, staying indoors.

Current warning systems for air pollution alerts are intended, in fact, to reduce exposures of susceptible persons in this fashion.

The second error term $\bar{x}_t - z_t^*$ is non-Berksonian and has the greatest potential to introduce bias in the estimate \hat{b}_z when z_t^* is correlated with $\bar{x}_t - z_t^*$. Even if the terms are uncorrelated so that \hat{b}_z will be a roughly unbiased estimate of b_x , it will reduce efficiency relative to a study in which \bar{x}_t is available since z_t and $\bar{x}_t - z_t$ share the same coefficient in (E8).

The difference $\bar{x}_t - z_t^*$ between average personal exposures and the true ambient value can be analyzed further by considering an individual personal exposure x_{it} .

Because individual i 's exposure on day t derives either from indoor or ambient sources, we can write $x_{it} = a_{it}z_t^* + (1 - a_{it})I_{it}$ where I_{it} is the concentration of pollutant generated by indoor sources such as tobacco smoke and pets, a_{it} is his or her fraction of exposure from ambient sources which take place either outdoors or result from penetration of ambient pollution indoors. It follows that $\bar{x}_t = \bar{a}_t z_t^* + \bar{I}_t$ where

$$\bar{I}_t = \sum_i (1 - a_{it})I_{it} / n_t.$$

That is, the average personal exposure is proportional to the

ambient level, offset by the effects of the population average of the non-ambient indoor sources.

Wilson and Suh (27) have argued that the daily population average concentrations of fine particles derived from indoor sources \bar{I}_t are approximately independent of ambient levels z_t across time. When this is true, failure to measure indoor sources will not introduce further bias in the estimation of b_x because the

deviations due to indoor air exposure are a second example of Berkson error, which errors will tend to cancel one another out when averaged over the population.

Nevertheless, z_t^* is only proportional to \bar{x}_t so that even if \bar{a}_t varied little over time

($a_t \approx a$), the coefficient \hat{b}_z from a regression of y_t on z_t^* would estimate ab_x , not b_x .

Hence, if 20% of daily exposure results from indoor sources independent of the ambient levels, the regression on ambient levels will yield coefficients that are roughly 20% smaller than would have occurred with actual personal exposures. However, this may be the appropriate coefficient for policy makers seeking an estimate of the effect of an inarguable measure of ambient levels. This, however, assumes that particles from indoor sources and outdoor sources are identical; that is, they are similar in composition and toxicity. If not the case, then the two types of particles are more appropriately treated as separate pollutants, and the personal exposure measure desired would be $a_{ii}z_t^*$, the personal exposure to particles from outdoor sources. Studies using sulfates as a tracer for particles from outdoor sources indicate that indoor/outdoor ratios are less than one. Since people spend most of their time indoors this suggests that a_{ii} will be less than one, and that the second term in Equation (E8) will be negatively correlated with z_t , and will bias the estimated coefficient downward. This also illustrates that the model is not restricted to cases where $E(x) = E(z)$.

The final of the three error terms in Equation E8, $z_{ii}^* - z_t$, represents the instrument measurement error in the ambient levels; like $x_{ii} - \bar{x}_t$; it is close to the Berkson type. This term would tend to be cancelled out by spatial averaging across multiple, unbiased ambient monitors in a region. For example, Kelsall et al. (26)

averaged daily TSP data from up to nine monitors in their analysis of effects of particles on mortality in Philadelphia. However, in many cities there is only one or a few monitors operating at a time. Even with a small number of monitors, longer-term drift in instruments will not substantially affect estimates of \mathbf{b}_x because the time series models control for such trends by inclusion of $s(t)$ in (E6). For this final error term to cause substantial bias in $\hat{\mathbf{b}}_z$, the error $z_t^* - z_t$ must be strongly correlated with z_t at shorter time scales. Further investigations of this correlation in cities with many monitors are warranted.

To summarize, we have discussed three components of measurement error: 1) an individual's deviation from the risk-weighted average personal exposure; 2) the difference between the average personal exposure and the true ambient level; and 3) the difference between the measured and the true ambient levels, which includes spatial variation and instrument error. Our analysis argues that the first and third components are of the Berkson type and, therefore, are likely to have smaller effects on the relative risk estimates. However, the second component can be a source of substantial bias, if for example, there are short-term associations of the contributions of indoor sources with ambient concentrations. In the following section, we present one simple analysis of the PTEAM data that illustrates how we can further study the effects of the most important second component.

4. Evaluating Potential Measurement Error Bias in Pollutant-Mortality Relative Risk Estimates

The framework set out in the previous section can be used, in combination with data on the components of error, to quantify the consequences of exposure measurement error. In this section, we use one of the few available data sets with ambient and personal measurements to illustrate one approach. In Section 4.1, we begin by using daily measurements of personal exposure for 178 persons followed in the PTEAM Study (23) to quantify the difference between concentration measured by an ambient monitor and the average of personal exposures. In Section 4.2, we present one approach for estimating the size of bias in estimated PM_{10} -mortality regression coefficients \hat{b}_z as an estimate of the true relative risk for personal exposure b_x , due to having data from one or a few ambient monitors rather than personal exposure data for PM_{10} .

4.1 PTEAM Study Data

The PTEAM Study (23, 28) generated a daily measurement of personal exposure to PM_{10} for a sample of 178 nonsmoking residents of Riverside, California aged 10 years or more for the period September 22 through November 9, 1990. In addition, a daily average PM_{10} value from an ambient monitor positioned near the homes was also collected; Pellizzari and Spengler provide details on the methods used to collect these data (29).

We use the PTEAM Study data to estimate the correlation between the daily PM_{10} concentration for the ambient monitor z_t and the difference between the average

personal exposure and concentration measured by the ambient monitor $\bar{x}_t - z_t$. These estimates correctly account for the varying number of observations on a given day. But the average personal exposure value is based on relatively few measurements and is therefore more variable across time than the actual mean exposure. Note that Equation (E8), includes a weighted average of personal exposures, with weights determined by the baseline personal risk for each individual. In the PTEAM Study, those weights are unavailable and hence, an unweighted average is used. **Figure 2** displays a time-series plot of the daily ambient values and the average personal exposures. The correlation across *time* of these two series is estimated to be 0.58 (95% confidence interval 0.35 to 0.74). We note that this is much greater than the more widely cited cross-sectional correlation from this study. It would likely be even greater if the mean personal exposure was calculated on a larger number of persons each day. The corresponding correlation across time between the ambient monitor concentrations and the daily differences between the personal and ambient values is -0.63 with 95% confidence interval, -0.77 to -0.42. Hence, the hypothesis that the measurement error $\bar{x}_t - z_t$ is uncorrelated with z_t is not consistent with the PTEAM Study data. Some bias in the regression coefficient is therefore expected. Because the correlation of $\bar{x}_t - z_t$ and z_t is negative, the coefficient \hat{b}_z in the regression on z_t will tend to *underestimate* the coefficient in the regression on \bar{x}_t in a single pollutant analysis. We now assess the size of the bias that will result from this measurement error.

4.2 Addressing the Bias in PM₁₀-Mortality Regression Coefficients

In this section, we illustrate how the PTEAM Study results of Sections 4.1 or other, perhaps more appropriate data sets on the difference between average risk-weighted personal exposure and ambient monitor concentrations, can be used to estimate bias in the results of log-linear regression models.

If available, we would have used the average personal exposure series, \bar{x}_t , for at risk residents of each city in the standard log-linear regression model rather than z_t , as was used in the original analyses. We would then compare the regression coefficients obtained when \bar{x}_t is the predictor with those using z_t to assess the bias.

Obviously, \bar{x}_t is not available except in special circumstances. But from the PTEAM Study data, shown in **Figure 2** or similar data, we can estimate the relationship of \bar{x}_t and z_t , for example, by assuming:

$$\bar{x}_t = \mathbf{q}_0 + \mathbf{q}_1 z_t + \mathbf{e}_t \quad (\text{E9})$$

where \mathbf{q}_0 and \mathbf{q}_1 are the intercept and slope to be estimated from the available data.

We can then use the fitted (E9) to predict the unobserved \bar{x}_t from the available z_t and

then use the predicted value $\hat{\bar{x}}_t$ as the desired exposure values when estimating the

pollution-mortality relative risk \mathbf{b}_x . In fact the estimate of \mathbf{b}_x has the simple form

$\hat{\mathbf{b}}_x = \hat{\mathbf{b}}_z / \hat{\mathbf{q}}_1$. This is one well-known approach to adjust for exposure measurement error

called “regression calibration” (7). As an illustration, we have applied this strategy to a regression of daily mortality on ambient concentrations of PM₁₀ for Riverside, California

for the period 1987-1994. We estimate $\hat{\mathbf{q}}_0 = 59.95$ (se = 7.21), $\hat{\mathbf{q}}_1 = 0.60$ (se = 0.080), and

var(ϵ) = 22.4.

Calibration is easy to implement and apply. Its limitations are that confidence intervals for $\hat{\mathbf{b}}_x$ depend upon large sample theory and it does not extend easily to situations where multiple sources of information about the \bar{x}_t, z_t relationship are available.

It is simple, however, to overcome these possible limitations of calibration by using a simulated value \bar{x}_t^* rather than the predicted value $\hat{\bar{x}}_t$ from (E9). That is, we use Equation (E9) to simulate the average personal exposure, \bar{x}_t^* , from the ambient exposure, z_t , for a city or period of interest when \bar{x}_t is not available, under the assumption that the estimated θ s and $\text{var}(\varepsilon)$ are applicable. This simulated series \bar{x}_t^* is then used instead of z_t in the log-linear regression. The result is one estimate of \mathbf{b}_x - call it $\hat{\mathbf{b}}_x$. If we then repeatedly simulate \bar{x}_t^* s and for each, fit the log-linear regression to obtain $\hat{\mathbf{b}}_x$ we obtain a distribution of $\hat{\mathbf{b}}_x$ s. The difference between the mean of the simulated $\hat{\mathbf{b}}_x$ s and the $\hat{\mathbf{b}}_z$ derived from the log-linear regression of mortality on \bar{z}_t^* , is a measure of the bias resulting from having z_t rather than the average personal exposure for that city. By simulating \bar{x}_t^* s rather than using a fixed predicted value $\hat{\bar{x}}_t$, we properly account for non-linearities and sources of variation in $\hat{\mathbf{b}}_x$ and can extend the analysis to more complicated situations.

Figure 3 shows the distribution of the $\hat{\mathbf{b}}_x$ s for Riverside (solid curve). Also shown is the normal approximation of the likelihood function for the coefficient $\hat{\mathbf{b}}_z$ from the log-linear regression of mortality directly on z_t (dotted curve). Solid and dotted lines

are at the centers of these distributions. We find that the $\hat{\mathbf{b}}_x$ s have a mean 1.42% increase in mortality (95% interval -0.11, 2.95) per 10 unit change in PM_{10} . In comparison, the estimate of \mathbf{b}_z from the usual log-linear model (dashed vertical line) is $\hat{\mathbf{b}}_z = 0.84\%$ (95% interval -0.06, 1.76). Hence, measurement error has biased the result toward the null. Second, the distribution of the $\hat{\mathbf{b}}_x$ s is more dispersed than the distribution of $\hat{\mathbf{b}}_z$. This is because we have taken into account the variability due to having z_t , not \bar{x}_t , i.e., arising from $\text{Var}(e_t)$ in Equation (E9). The results are very similar to what we obtain from calibration.

This calculation assumes the estimated relationship between \bar{x}_t and z_t for the PTEAM Study is the true one, and hence, we ignore a second component of uncertainty due to estimation of the relationship between \bar{x}_t and \bar{z}_t from the finite sample size of the PTEAM Study data taken at one site and a particular time period. That is, even if we assume that the relationship between \bar{x}_t and z_t is known, estimating the association of mortality with \bar{x}_t is less precise than with z_t , given only z_t in that particular city. Of course, the relationship of \bar{x}_t and z_t is not precisely known and needs to be quantified further. Dominici et al. (30) provide a more complete analysis of the bias in $\hat{\mathbf{b}}_z$ as an estimate of \mathbf{b}_x using the PTEAM Study and four other data sets and a more complete statistical model. Their findings are qualitatively similar to those presented here. Finally, it is important to note that our assessment of bias assumes that the health effects of personal exposure to particles originating outdoors and indoors are the same. To assume otherwise would require substantially more detailed data and modeling.

5. Summary and Research Recommendations

The differences between true personal exposure for every individual (x_{it}) and measured ambient concentrations, averaged over a few fixed, imprecise monitors (z_t), is inherently complex, as is the effect of this exposure measurement error on estimates of pollution-mortality relative risks. Nonetheless, it is useful and imperative to analyze these effects in light of our current understanding of the measurement process. This paper presents one framework for doing so. We distinguish two extremes of a continuum of types of measurement errors: Berkson and classical errors. The former is likely to create little bias in mortality-relative risk estimates; the latter has more serious consequences.

We posit a relative risk model in which an individual's hazard of death on a given day is expressed as a function of his or her personal exposure which is decomposed to highlight three types of exposure errors. This model is then aggregated to produce the model for the expected total deaths in a population used in most time-series analyses. This model shows that a risk-weighted average personal exposure measure is the desired one and we discuss the consequences of the widely used feasible alternative, ambient concentration. In contrast, differences between individual exposures on a given day and the risk-weighted average of personal exposures are an example of Berkson error and not likely to cause substantial bias in coefficients from time-series morbidity studies. Our analysis suggests that the largest biases in inferences about the mortality-personal exposure relative risk will occur due to the more complex errors between ambient and average personal exposure measures. If indoor sources produce particles of similar composition and toxicity as outdoor source particles, indoor sources

may be a major component of this error. Finally, as an illustration we have used the best available data, that from the PTEAM Study in Riverside, California, with both personal exposure and ambient time series to quantify the size of this error. Our analysis indicates that the coefficient obtained from regressing mortality on measured ambient levels (z_t) is smaller than what we expect if we regress mortality on average personal exposure (\bar{x}_t).

For tractability and clarity, we have conducted a first-order analysis of exposure errors and have ignored possible second and higher order effects in which daily fluctuations in the variance of personal exposures across a population or in the covariations among the measurement errors could introduce additional biases. Second order terms will be insignificant in studies of particulate effects on mortality where the first order terms are on the order of percent. Such higher-order analyses for other studies of for example morbidity, are beyond the scope of this paper and will require substantially more detailed models and data. It is, however, possible that higher order effects are important so that further investigation is necessary.

Epidemiologic research is necessarily limited by the quality of the health outcome and risk factor measurements (31). Time-series studies of the acute effects of air quality on mortality are subject to the limitations posed by the available measurements of pollution levels. The generic criticism -- that measurement errors render the results of such time-series models uninterpretable -- is incorrect. This paper demonstrates that the consequences of measurement error can be quantified, although only a few informative data sets are presently available. This paper suggests that differences between the average personal exposure and ambient measurements are the most likely

source of substantial bias. We suggest that data should be collected for comparison of risk-weighted average personal exposure with ambient levels in several cities with varying degrees of spatial heterogeneity in ambient levels, population composition, and indoor pollution sources. Given such data, models like those summarized by Dominici et al. (32) can be used to quantify more precisely the biases due to pollutant measurement errors.

This paper focuses on the effects on relative risk estimates of using z_t , measured ambient particle levels rather than x_{it} , actual personal exposures in log-linear regressions. Such effects are important from a scientific perspective to quantify the health risks of exposure to particulate pollution. From a regulatory perspective, the effect of having the imprecise z_t rather than the “true” ambient value z_t^* may be of greater interest since it is ambient levels that may or may not be regulated further. A more detailed error analysis of the $z_t - z_t^*$ difference would investigate the spatial variation in particulate levels and how the number of monitors used to calculate z_t reduced this source of measurement error.

The analyses in Sections 3 and 4 focus on measurement error in a single pollutant measure, PM_{10} . As discussed in Section 2, simultaneous errors in several pollutants can complicate the analysis. Section 2 clearly demonstrates, however, that qualitative biases – that is, changes in the sign of a coefficient – can occur only when the measurement errors for different pollutants are highly correlated with one another. This level of correlation might arise if two or more pollutants are measured by the same instrument (e.g., different fractions of PM) or if multiple instruments are housed in the same location which is subject to atypical exposure patterns. The possibility

nevertheless requires detailed investigation, since in this case the findings of epidemiologic studies could be misleading. Personal exposure studies that collect multiple exposures can provide the necessary data to investigate the effects of co-occurring errors using straightforward extensions of the approaches outlined in Section 3 and 4.

In this paper, we have considered the effects of exposure measurement error on regression coefficients from log-linear models in which serial correlation is accounted for using flexible smoothing splines. An alternate analytic strategy is to fit a linear regression with time series errors (ARIMA model, (33)). In certain specific time series models, the degree of attenuation due to classical error might be reduced since to account for the autocorrelated errors, the ARIMA filters or smooths both the responses and the predictors which might reduce the degree of measurement error. Further research on this possibility is warranted.

The measurement error framework in Section 3 and the illustrative calculations in Section 4 make apparent several open questions and opportunities for additional data collection that would enable more accurate quantification of the effects of measurement error in assessing the air pollution-mortality relationship. In relation to single-pollutant models, we consider that the two most important questions are:

- Is the average personal exposure to pollutants from indoor sources correlated over time with ambient levels?
- Does the difference between baseline risk-weighted average exposure and population average exposure vary slowly over time?

For models with multiple pollutants, the additional key question is:

- How do the components of error identified in Equation (E5) co-vary across pollutants? For example, how do the differences between actual ambient levels and the measured levels correlate across the different pollutants and how do these differences depend on the true values of other pollutants or covariates?

Wilson and Suh (27) have conducted a meta-analysis of data from multiple sites and conclude, in answer to the first question above, that concentrations of fine particles originating from indoor sources are independent of ambient levels over time. To confirm this finding and to address the remaining, key questions, additional research is warranted. It would be highly informative if, in several cities with diverse pollution sources and patterns, a stratified sample of the population were drawn with one stratum representing the entire population and the second representing the frail subgroup. Daily measurements of personal exposure and indicators of indoor sources would be collected for multiple pollutants for each person. Ambient levels would also be monitored. Decisions about the numbers of persons within each subgroup and the numbers of days of monitoring for each person would be made based upon preliminary analyses of data from one city.

Table 1: Predicted Bias in Bivariate Regression Coefficients under Different Covariance Structures for True Exposures and Measurement Errors When Both Variables Have a True Effect: $b_{x_1} = b_{x_2} = 1.0$. We assume $\text{Var}(x_1) = \text{Var}(x_2) = 1$.

$\text{Corr}(x_1, x_2)$	$\text{Var}(d_1)$	$\text{Var}(d_2)$	$\text{Corr}(d_1, d_2)$	$E(\hat{b}_{z_1})$	$E(\hat{b}_{z_2})$
0.0	1.0	1.0	0.0	0.50	0.50
0.5	1.0	1.0	0.0	0.60	0.60
-0.5	1.0	1.0	0.0	0.33	0.33
0.0	1.0	1.0	0.5	0.40	0.40
0.0	1.0	1.0	-0.5	0.67	0.67
0.0	0.5	2.0	0.0	0.67	0.33
0.5	0.5	2.0	0.0	0.71	0.53
0.5	0.5	2.0	0.3	0.66	0.27
0.5	0.5	2.0	0.5	0.64	0.21
0.5	0.5	2.0	0.7	0.64	0.14
0.5	0.5	2.0	-0.5	0.83	0.50
0.5	0.5	2.0	-0.7	0.91	0.57
0.5	0.5	2.0	-0.9	1.00	0.66

Table 2: Predicted Bias in Bivariate Regression Coefficients under Different covariance Structures for True Exposures and Measurement Errors
When Only One Variable Has a True Effect: $b_{x_1} = 0, b_{x_2} = 1$. We assume $\text{Var}(x_1) = \text{Var}(x_2) = 1$.

$\text{Corr}(x_1, x_2)$	$\text{Var}(d_1)$	$\text{Var}(d_2)$	$\text{Corr}(d_1, d_2)$	$E(\hat{b}_{z_1})$	$E(\hat{b}_{z_2})$
0.0	0.5	2.0	0.0	0.00	0.33
0.0	0.5	2.0	0.5	-0.12	0.35
0.0	0.5	2.0	-0.5	0.12	0.35
0.5	0.5	2.0	0.0	0.06	0.29
-0.05	0.5	2.0	0.0	-0.06	0.29
0.5	0.5	2.0	0.3	-0.01	0.28
0.5	0.5	2.0	0.5	-0.07	0.29
0.5	0.5	2.0	0.7	-0.15	0.29
0.5	0.5	2.0	-0.5	0.17	0.33
0.5	0.5	2.0	-0.7	0.21	0.36
0.5	0.5	2.0	-0.9	0.26	0.39

References

1. Armstrong BK, Saracci R, White E. Principles of exposure measurement in epidemiology. New York:Oxford University Press, 1992.
2. Thomas D, Stram D, Dwyer J. Exposure measurement error: Influence on exposure-disease relationships and methods of correction. *Annu Rev Public Health* 14:69-93 (1993).
3. National Research Council (NRC), Committee on Advances in Assessing Human Exposure to Airborne Pollutants. Human Exposure Assessment for Airborne Pollutants: Advances and Opportunities. Washington, D.C.:National Academy Press, 1991.
4. Ferris BG, Jr., Speizer FE, Spengler JD, Dockery DW, Bishop YM, Wolfson M, Humble C. Effects of sulfur oxides and respirable particles on human health: Methodology and demography of populations in study. *Am Rev Respir Dis* 120:767-779 (1979).
5. Dockery DW, Pope CA, III, Xu X, Spengler JD, Ware JH, Fay ME, Ferris BG, Jr., Speizer FE. An association between air pollution and mortality in six U.S. cities. *N Engl J Med* 329:1753-1759 (1993).
6. Fuller WA. Measurement Error Models. New York, New York:John Wiley & Sons, 1987.
7. Carroll RJ, Ruppert D, Stefanski LA. Measurement error in nonlinear models. London:Chapman and Hall, 1995.

8. Spiegelman D, McDermott A, Rossner B. Regression calibration method for correcting measurement-error bias in nutritional epidemiology. *Am J Clin Nutr* 65:1179S-1186S (1997).
9. Willett W. Correction for the effects of measurement error. In: *Nutritional Epidemiology* (Willett W, ed). New York:Oxford University Press, 1998;301-320.
10. Pierce DA, Stram DO, Vaeth M. Allowing for random errors in radiation dose estimates for the atomic bomb survivor data. *Radiat Res* 123:275-284 (1990).
11. Shy CM, Kleinbaum DG, Morgenstern H. The effect of misclassification of exposure status in epidemiological studies of air pollution health effects. *Bull N Y Acad Med* 54:1155-1165 (1978).
12. Goldstein IF, Landovitz L. Analysis of air pollution patterns in New York City -- I. Can one station represent the large metropolitan area? *Atmos Environ* 11:47-52 (1977).
13. Goldstein IF, Landovitz L. Analysis of air pollution patterns in New York City -- II. Can one aerometric station represent the area surrounding it? *Atmos Environ* 11:53-57 (1977).
14. Navidi W, Thomas D, Stram D, Peters J. Design and analysis of multilevel analytic studies with applications to a study of air pollution. *Environmental Health Perspectives* 102:Suppl-32 (1994).
15. National Research Council (NRC), Commission on Life Sciences, Board on Toxicology and Environmental Health Hazards, Committee on the Epidemiology of Air

Pollutants. Epidemiology and Air Pollution. Washington, D.C.:National Academy Press, 1985.

16. Dockery DW, Pope CA, III. Acute respiratory effects of particulate air pollution. *Annu Rev Public Health* 15:107-132 (1994).

17. Lipfert FW, Wyzga RE. Air pollution and mortality: the implications of uncertainties in regression modeling and exposure measurement. *J Air Waste Manage Assoc* 47:517-523 (1997).

18. Lipfert F. Clean air skepticism. *Science* 278:19-20 (1997).

19. Lipfert FW. Air pollution and human health: perspectives for the '90s and beyond. *Risk Anal* 17:137-146 (1997).

20. Schwartz J, Dockery DW, Neas LM. Is daily mortality associated specifically with fine particles? *J Air & Waste Manage Assoc* 46:927-939 (1996).

21. Wallace L. Indoor particles: A review. *J Air Waste Manage Assoc* 46:98-126 (1996).

22. Janssen NA, Hoek G, Brunekreef B, Harssema H, Mensink I, Zuidhof A. Personal sampling of particles in adults: relation among personal, indoor, and outdoor air concentrations. *Am J Epidemiol* 147:537-547 (1998).

23. Ozkaynak H, Xue J, Spengler J, Wallace L, Pellizzari E, Jenkins P. Personal exposure to airborne particles and metals: results from the Particle TEAM Study in Riverside, California. *J Expo Anal Environ Epidemiol* 6:57-78 (1996).

24. Snedecor GW, Cochran WG. Statistical Methods. Ames, Iowa: The Iowa State University Press, 1980.

25. Carroll RJ, Spiegelman CH, Lan KKG, Bailey KT, Abbott RD. On errors in variables for binary regression models. *Biometrika* 71:19-25 (1984).

26. Kelsall JE, Samet JM, Zeger SL, Xu J. Air pollution and mortality in Philadelphia, 1974-1988. *Am J Epidemiol* 146:750-762 (1997).

27. Wilson WE, Suh HH. Fine particles and coarse particles: concentration relationships relevant to epidemiologic studies. *J Air Waste Manage Assoc* 47:1238-1249 (1997).

28. Mendelsohn R, Orcutt G. An empirical analysis of air pollution dose-response curves. *J Environ Economics Manage* 6:85-106 (1979).

29. Pellizzari E, Spengler J. Particle Total Exposure Assessment Methodology (PTEAM): Pilot Study, Volume II: Protocols for Environmental Sampling and Analysis. Work Plan for EPA Contract No. 68-02-4544, EPA Work Assignment 67, CARB Agreement No. A833-060. Research Triangle Park: US Environmental Protection Agency, 1990.

30. Dominici F, Samet J, Xu J, Zeger S. Combining evidence on air pollution and daily mortality from the largest 20 U.S. cities: a hierarchical modeling strategy. *J Royal Stat Soc, Series C* (1999).

31. Vedal S. Ambient particles and health: lines that divide. *J Air Waste Manag Assoc* 47:551-581 (1997).

32. Dominici F, Zeger S, Samet J. A measurement error correction model for time-series studies of air pollution and mortality. Technical Report. Baltimore, MD:Johns Hopkins University, 1999.

33. Fuller WA. Introduction to statistical time series. New York:Wiley & Sons, 1996.