

A MEASUREMENT ERROR MODEL FOR BINARY AND ORDINAL REGRESSION

Daniel W. Schafer

Department of Statistics, Oregon State University

Corvallis, Oregon 97331-4606

Leonard A. Stefanski

Department of Statistics, North Carolina State University

Raleigh, North Carolina 27695-8203

Tor D. Tosteson

Channing Laboratory, Harvard Medical School

Boston, Massachusetts 02115-5899

TITLE: A Measurement Error Model for Binary and Ordinal Regression

SUMMARY

A recent study of nitrogen dioxide exposure and respiratory health makes use of improved exposure assessment using indoor monitoring. This paper develops and applies methods for estimating the parameters of binary and ordinal regression models when the measured indoor concentrations are viewed as surrogates for personal exposure. Data from two independent studies are used to provide information on the relationship between measured indoor concentrations and personal exposure. Methods for evaluating the assumptions necessary to estimate the parameters of our models are discussed.

KEY WORDS binary and ordinal regression probit models errors-in-variables.

INTRODUCTION

In this paper, we develop and apply a measurement error model for binary and ordinal regression to data relating the occurrence of respiratory symptoms in children to indoor concentrations of nitrogen dioxide (NO_2). These data were collected in Watertown, Massachusetts during the course of a Harvard study¹ designed to evaluate the role of common air pollutants in the development of lung disease. Originally this study only collected outdoor pollutant concentrations, but in light of emerging evidence that some pollutants such as NO_2 have important indoor sources, and given the fact that children spend a large portion of their time indoors, the investigators decided to place monitors directly in the houses of a subsample of their original population.

In our analysis, we adopt the point of view that the measured indoor concentrations are surrogates for a child's true exposure to NO_2 , and that this "personal" exposure is actually the predictor variable of interest. We are able to determine the relationship between personal exposure and indoor concentrations by using data from two studies employing both personal monitoring and indoor monitoring.^{2,3} If we then assume that indoor concentrations only affect respiratory health through their contribution to personal exposure, we can use a measurement error method to estimate the association between personal exposure and respiratory symptoms.

Because our primary goal is the development of an appropriate methodology, we have chosen to analyze only a greatly simplified subset of the data from the constituent studies. In particular, we have omitted covariates known to be important determinants of respiratory health, although we have made our models flexible enough to include these variables in subsequent applications. In addition to limiting the scientific validity

of our results, the omission of relevant covariates may violate the conditional independence assumption we use to derive our measurement error model. We discuss methods for testing the appropriateness of this assumption, and formulate a somewhat weaker assumption which may be more appropriate in this case, although for the more complete data set for which these methods are ultimately intended the stronger version seems justified.

DESCRIPTION OF THE DATA SETS

The data set from Watertown includes two health variables and two indoor measurements for each of 231 children. The health variables are derived from a questionnaire administered in the summer of 1984 to the parents of the children, and include binary responses to questions concerning the occurrence of any wheezing apart from colds in the preceding 12 months, and presence of persistent wheeze in the preceding 12 months. The second symptom, persistent wheeze, is considered a clinically more significant outcome, but the relatively low prevalence of this symptom prompted the inclusion of the first indicator, which we will sometimes refer to simply as wheeze.

The indoor concentrations of NO_2 were measured by placing passive diffusion tubes⁴ in the bedroom and kitchen of the children's homes. The tubes were left in place for two seven-day periods during the winter of 1984-1985. We only use the measurements from the first period in this analysis.

The two studies relating personal exposure to indoor concentrations were conducted in Portage, Wisconsin, during the winter of 1981-1982 and in the Netherlands during the winter of 1984-1985. Quackenboss et al.² present an analysis of the Portage data with a particular emphasis on the effect of

daily activity patterns. Houthuijs et al.³ present results from the Dutch data set on the effects of NO_2 and tobacco smoke exposure on the respiratory health of children.

The data set derived from the Portage study consists of three variables for 81 children living in separate homes. One variable represents a measurement of personal exposure over a seven-day period, obtained by passive diffusion tubes worn on the lapels of the children. The other two variables are seven-day samples of kitchen and bedroom concentrations obtained during the same week, also by means of the diffusion tube.

The data set from the Dutch study contains the same three variables for 564 children living in five small, non-industrial communities in the southeastern part of the country. As in the Portage and Watertown studies, the children are from separate homes.

A MODEL FOR THE DATA

We propose a model that addresses the two unique features of our data. The first is the presence of two closely related binary outcomes. As it happens, the occurrence of persistent wheeze implies the occurrence of any wheeze apart from colds. We take advantage of this relationship between the outcomes by combining them into a single ordinal outcome with three categories: no wheezing, some wheezing apart from colds, and persistent wheezing. We then develop a method applicable to both ordinal and binary regression models.

The second feature of the data is the availability of two surrogates for the predictor variable of interest, namely the kitchen and bedroom concentrations of NO_2 . We accommodate this feature by using the Portage and Dutch data sets to develop a predictive model for personal exposure given

kitchen and bedroom concentrations. From this model we obtain predicted values of personal exposure for each child in the Watertown data.

Recognizing that these predicted values can be represented as the sum of the true exposures plus an additive error term, we then construct an appropriate errors-in-variables model for modeling the dependence of wheeze on personal exposure. This approach can be defended only if the predictive model developed from the Dutch and Portage data is portable to the Watertown data, and the errors-in-variables model employed recognizes that the error arises not from an independent measurement of personal exposure but rather from use of a predictive model.

We start by postulating an underlying continuous latent variable measuring proneness to wheeze, Y^* . For example, for a binary model of persistent wheeze (PW), it is assumed that $PW = 0$ whenever $Y^* \leq c_1$ and that $PW = 1$ for $Y^* > c_1$. For the model which regards wheeze (W) and persistent wheeze as ordered categories, two cut points are introduced; $Y^* \leq c_1$ corresponds to $(W, PW) = (0, 0)$, $c_1 < Y^* \leq c_2$ to $(W, PW) = (1, 0)$, and $Y^* > c_2$ to $(W, PW) = (1, 1)$.

Now assume that in each of the three populations for which the Dutch, Portage and Watertown data sets are representative, Y^* has a joint normal distribution with $X_1 = \log$ personal exposure (PE), conditionally on a vector of covariates X_2 , and a vector Z containing the log kitchen and bedroom concentrations (KC & BC). The assumption of conditional normality leads to a tractable form for the likelihood of the observed data, and is empirically justified in our example. However, the models can be developed under distributional assumptions much more general than those of joint normality, and are likely to be at least approximately correct under still weaker conditions. We introduce X_2 for added generality, noting that in the

more complete data set for which these methods are eventually intended, adjustments will have to be made for important covariates. We will sometimes write X for $(X_1', X_2')'$.

Recall that the Dutch and Portage data sets contain observations on only X_1 , X_2 , KC and BC whereas the Watertown data contains observations on only Y^* (actually the discretized version of Y^*), X_2 , KC, and BC. Because none of our data sets contain simultaneous observations on Y^* and X_1 , it is not possible to directly estimate $\text{Cov}(Y^*, X_1 | X_2, Z)$. Thus even under joint normality for $(Y^*, X_1 | X_2, Z)$ our goal of investigating the association between personal exposure to NO_2 and respiratory health cannot be realized unless some restrictions are imposed on $\text{Cov}(Y^*, X_1 | X_2, Z)$. The most natural restrictions in our setting involve the role of KC and BC in determining respiratory health. An assumption which permits us to estimate the parameters of the distribution of $Y^* | X$ can be constructed by (i) asserting that indoor NO_2 concentrations can only affect the respiratory system through their effect on personal exposure to NO_2 , and (ii) arguing that KC and BC carry no additional information relevant to health status that is not already contained in the covariate vector X_2 . For example, the fact that KC and BC are both concentrations suggests a possible relationship to house size and hence to numerous other socio-economic variables which may be related to health status. The completeness argument, (ii), claims that all such covariable information is contained in X_2 . A logical consequence of (i) and (ii) is that the distributions of $Y^* | X$ and $Y^* | (X, Z)$ are equal, which in turn allows us to estimate the parameters of $Y^* | X$ from the observed data. For reference we state this as an assumption.

(C1) The conditional distributions of $Y^* | X$ and $Y^* | (X, Z)$ are equal.

An alternative assumption that also allows the estimation of the parameters of the distribution of $Y^*|X$ is:

(C2) The conditional distributions of $Y^*|X$ and $Y^*|(X, E(X_1|X_2, Z))$ are equal.

Assumption (C2) is technically weaker than (C1) in that (C1) implies (C2) but not vice versa. If one adopts the view that interest lies solely in the distribution of $Y^*|X$, irrespective of the role played by Z , (C2) may be more appropriate than (C1). In the more complete data set for which our methods are intended, it can be argued that (C1) has a physiological basis.

Under our normality assumption, (C2) is equivalent to $0 = \text{Cov}(Y^*, X_1 - E(X_1|X_2, Z))^2$ which is equal to $\text{Var}(Y^*)\text{Var}(X_1 - E(X_1|X_2, Z))\text{Corr}(Y^*, X_1 - E(X_1|X_2, Z))^2$, and thus is at least approximately valid when the prediction error in $E(X_1|X_2, Z)$ is small or the correlation of Y^* with the error of X_1 about its regression on X_2 and Z is small. While the latter is difficult to verify, the former can easily be checked. However, the objective of our study is to develop models for situations in which the prediction error is not negligible.

Operationally, (C1) and (C2) are nearly identical, in that the same model can be derived from either. The distinction between the two arises only in the interpretation of a test statistic for model adequacy which we present later in this section.

We are now in a position to proceed with the development of a model. Separate analyses of the Dutch and Portage data (see the next section) showed remarkable agreement between the two linear regression models of X_1 on (X_2, Z) and provided ample evidence for combining the data to obtain a

final prediction model for X_1 . Because the combined data contain over three times as many observations as the Watertown data, the variances of prediction are essentially equal to the MSE over much of the range of interest, so there is some heuristic support for ignoring the sampling variability in a consistent estimate of the coefficients. We do this in part to simplify the discussion and partly because it seems reasonable in the present context. However, in other problems in which the discrepancy between the sizes of the validation and target data sets is not so great, it may be necessary to incorporate this additional source of variation into the model.

To use the regression of X_1 on (X_2, Z) from the Dutch-Portage population to predict X_1 values for the Watertown data, it is necessary to assume that the regressions in the two populations coincide. Letting μ_1 and σ_1^2 denote the (known) regression function and prediction-error variance from the Dutch-Portage population we assume:

(M1) Within the Watertown population $X_1 | (X_2, Z)$ is distributed $N(\mu_1(X_2, Z), \sigma_1^2)$ where $\mu_1(\cdot, \cdot)$ and σ_1^2 are known.

We now define the variable $\bar{X}_1 = \mu_1(X_2, Z)$ which will play the role of a proxy for X_1 .

If we now assume

(M2) $Y^* | (X_1, X_2) \sim N(\beta_1 X_1 + \beta_2' X_2, \sigma^2)$

it follows from either (C1) or (C2) that

$$Y^* | (\bar{X}_1, X_2) \sim N(\beta_1 \bar{X}_1 + \beta_2' X_2, \sigma^2 + (\sigma_1 \beta_1)^2). \quad (1)$$

The vector X_2 can include a constant term.

Since our interest lies primarily in inference on β_1 , eqn. (1) indicates that an ordinary least-squares analysis of Y^* on (\bar{X}_1, X_2) is sufficient; the estimated coefficient of \bar{X}_1 will be unbiased for β_1 and the corresponding hypothesis test of $\beta_1 = 0$ will have the correct level although be less powerful than if X_1 was available rather than its proxy. This situation is similar to one which arises in the so-called Berkson (linear) errors-in-variables model. Unfortunately we do not observe Y^* but rather a discretized version of it and this introduces additional complications due to use of the proxy, \bar{X}_1 .

If we now postulate that the observed categorical variable, Y , falls into category i , $0 \leq i < k$, according to $c_i < Y^* < c_{i+1}$ where $-\infty = c_0 < c_1 < \dots < c_k = +\infty$, it follows from (M2) that

$$(M2') \Pr(Y = i) = \Phi[(c_{i+1} - (\beta_1 X_1 + \beta_2' X_2))/\sigma] \\ - \Phi[(c_i - (\beta_1 X_1 + \beta_2' X_2))/\sigma], \quad 0 \leq i < k.$$

That is, the induced model for $Y | (X_1, X_2)$ is a regression model for ordinal data of the type studied by McCullagh⁵ with a probit link. It is apparent that σ is not identified and there is no loss in generality setting $\sigma = 1$.

The induced model for $Y | (\bar{X}_1, X_2)$ becomes

$$\Pr(Y = i) = \Phi[(c_{i+1} - (\beta_1 \bar{X}_1 + \beta_2' X_2)) / (1 + (\sigma_1 \beta_1)^2)^{1/2}] \\ - \Phi[(c_i - (\beta_1 \bar{X}_1 + \beta_2' X_2)) / (1 + (\sigma_1 \beta_1)^2)^{1/2}], \quad 0 \leq i < k. \quad (2)$$

Thus the probit models for ordered categorical responses for both $(Y, (X_1, X_2))$ and $(Y, (\bar{X}_1, X_2))$ differ in that the coefficients of the latter model are attenuated by the factor

$$\gamma = (1 + (\sigma_1 \beta_1)^2)^{-1/2}.$$

Equations (M2') and (2) indicate that if $\hat{\theta}_1$ is the estimated coefficient of \bar{X}_1 in a probit model fit to $(Y, (\bar{X}_1, X_2))$ then a consistent estimate of β_1 in (M2') can be obtained according to

$$\hat{\beta}_1 = \hat{\theta}_1 (1 - (\sigma_1 \hat{\theta}_1)^2)^{-1/2} \quad (3)$$

provided that $(\sigma_1 \hat{\theta}_1)^2 < 1$; similar corrections can be made to the other parameters. Burr⁶ has studied estimation in models like (2). In particular, she suggests that confidence intervals for $\hat{\beta}_1$ be constructed by applying the transformation (3) to the limits of the confidence interval for $\hat{\theta}_1$.

Although our model was obtained under assumptions of joint normality a close inspection reveals that the same results are obtained under either assumption (C1) or (C2) provided: the probit link model in (M2) is appropriate; the regression structures in both the target and validation populations coincide with respect to both mean function (which need not be linear) and variance (which must be constant); and the prediction errors in the target population are normal. Clearly the most critical of these assumptions are (C1) (or (C2)) and the equality of regression structures, and we now discuss the consequences should they be violated. We consider only the case of linear prediction equations.

The simplest violation of our assumptions occurs when the prediction error variances are not equal. The consequence is that we misjudge the extent of attenuation. Understating the prediction error variance in the target population results in under estimation of γ and vice versa. There seems to be no way of checking this with the available data. In some studies it may be that the validation and target populations coincide in which case this would not be a problem.

Note that model (2) depends on Z through a particular linear combination dictated by $\mu_1(\cdot, \cdot)$ (see (M1)). If this regression plane is not appropriate in the target population then the derived model, (2), would likely depend on some other linear combination of KC , BC , and X_2 . Furthermore, even if we are willing to accept (M1) entirely, a violation of either (C1) or (C2) would have the effect that the induced model, (2), would not necessarily depend on X_2 and Z only through $\mu_1(\cdot, \cdot)$, and again we would expect some other linear combination to appear in (2). Furthermore there is a converse to this statement with regards to (C1). If we find that some other linear combination of KC and BC is more appropriate in (2) then it follows that (C1) is violated. However, this does not imply that (C2) is violated.

The variables KC and BC enter our derived model only through the linear combination dictated by (M1). The discussion of the critical assumptions (M1), (C1) and (C2) indicates that a violation of any one assumption generally has the effect of altering this linear combination. Consequently, we can define a test statistic which is sensitive to certain violations of assumptions, namely the likelihood-ratio statistic comparing the fits of the constrained and unconstrained models of Y on (\bar{X}_1, X_2) and Y on (Z, X_2) respectively. A small value of the test statistic indicates that the data

provide little evidence contradicting the assumptions. A large value of the statistic is more difficult to interpret. It cannot distinguish between a violation of (C1), (C2) or (M1); and in the case that only the distribution of $Y|X$ is of interest (and (C2) is assumed) it may merely reflect the fact KC and BC have some explanatory power in addition to PE and hence does not necessarily imply that (C2) is violated. An additional complication arises whenever the sampling variability in the estimate of $\mu_1(\cdot, \cdot)$ is appreciable, for then the statistic may be nonnegligible even though all assumptions are satisfied. Although the problems with this statistic are legion, it may be the only vehicle for checking assumptions in many situations and it seems like good statistical practice to compute it, provided it is interpreted in light of all its shortcomings.

RESULTS

An initial graphical analysis of the Dutch and Portage data showed that a linear fit for PE on KC and BC appeared reasonable and yielded normal residuals. To illustrate this point, Figure 1 shows a plot of PE versus KC in the Dutch data. Table 1 shows the coefficients for the fitted regression planes in the Dutch, Portage, and combined data sets. The results for the two original data sets are remarkably similar, and although a test for the equality of that residual variances is somewhat significant ($p \approx 0.06$), a formal test for the equality of the regression planes indicates no significant differences ($p \approx 0.76$) We therefore decided to use the combined data to form \tilde{X}_1 and to estimate σ_1^2 .

INSERT TABLE 1 HERE

INSERT FIGURE 1 HERE

Our initial analysis of the Watertown data focussed on the occurrence of wheeze, which the reader will recall is implied by the occurrence of persistent wheeze. Figure 2 is a scatterplot of BC versus KC for the 231 Watertown children. At each point is plotted a 1 or 0 corresponding to whether or not the child reported wheezing. While there is no apparent indication that wheeze prevalence increases with an increasing total of KC and BC, there is some evidence that there are more 1's in the northwest part of the point cloud than in the southeast. The line drawn on the points is the regression of BC on KC and is used, somewhat arbitrarily to show this effect. That is, above this line there is a higher percentage of children who wheeze .

INSERT FIGURE 2 HERE

This apparent effect is confirmed by a probit regression analysis. The fitted model for the proportion of wheezers in terms of the two predictor variables is the following:

$$\Phi^{-1}\{\text{Pr}(W=1)\} = -0.79 - 0.59 \times \text{KC} + 0.73 \times \text{BC} , \quad (4)$$

(0.32) (0.36)

where the numbers in parentheses are approximate standard errors. A likelihood ratio test for the joint significance of both predictors is suggestive ($p \approx 0.13$), although neither predictor is significant by itself. It can be seen that the ratio of the coefficients of KC and BC is suggestively close to -1, so that the effect of the two predictors may be

through their difference, or equivalently through the ratio of the untransformed concentrations.

Using values for \bar{X}_1 calculated from the coefficients reported for the combined data in Table 1, we obtain the following fitted model relating the occurrence of wheeze directly to the estimated PE's:

$$\Phi^{-1}(\Pr(W=1)) = -0.88 + 0.08 \times \bar{X}_1. \quad (5)$$

(0.21)

A 95% confidence interval for the slope is (-0.32, 0.48). It is interesting to note that corrections based upon equation (3) with σ_1^2 taken from the combined data produce no changes in these estimates for the number of digits reported. Thus, due to the fact that $\hat{\theta}_1$ and σ_1^2 are both relatively small, $\gamma \approx 1$, and consequently $\hat{\theta}_1$ and $\hat{\beta}_1$ are approximately equal in these data.

When we consider the ordinal variable, Y, we obtain similar results. A maximum likelihood fit of McCullagh's⁵ model for ordinal data with a probit link to Y and (KC, BC) results in the following estimated probabilities for the occurrence of wheeze ($Y > 0$) and persistent wheeze ($Y > 1$):

$$\Phi^{-1}(\Pr(Y>0)) = -0.73 - 0.51 \times KC + 0.62 \times BC.$$

(0.30) (0.34)

$$\Phi^{-1}(\Pr(Y>1)) = \Phi^{-1}(\Pr(Y>0)) - 0.93 \quad (6)$$

We note that $Y > 0$ is equivalent to $W = 1$; therefore the resemblance to the previously fitted model is not surprising. The fitted model using \bar{X}_1 is

$$\begin{aligned}\Phi^{-1}(\Pr(Y>0)) &= -0.79 + 0.05\bar{X}_1 \\ &\quad (0.20) \\ \Phi^{-1}(\Pr(Y>1)) &= \Phi^{-1}(\Pr(Y>0)) - 0.88\end{aligned}\tag{7}$$

As before, because $\gamma \approx 1$, the confidence intervals for the coefficients of X_1 and \bar{X}_1 are essentially identical.

The likelihood ratio test for the joint significance of KC and BC based upon the fitted ordinal regression (6) is less suggestive ($p \approx 0.20$) than the test based on the binary regression, and neither concentration is significant alone. This finding is due to the fact that the binary regression fit to PW, the results for which are not shown here, indicates a much weaker effect for KC and BC. Considering PW and W jointly in an ordinal regression model yields a p-value somewhere in between the p-values for the two binary regressions.

The difference in the deviances for the ordinal regressions (6) and (7) is 3.2, which when referred to the $\chi^2_{(1)}$ distribution yields a p-value of 0.07. The same difference calculated for the binary regressions (4) and (5) is 3.9, which yields a p-value of 0.05. This is the test statistic for checking the assumptions discussed in the last section. The p-values seem to indicate some evidence that the best explanatory model treats KC and BC separately; however, interpreting this observation is problematic. The most likely explanation is the absence of important covariates in our example data set, although as indicated earlier, other explanations exist. For instance, it might be the case that NO_2 affects respiratory health through some measure of exposure other than, or in addition to, the week-long integrated average used in the Dutch and Portage studies, and that this measure is correlated with kitchen and bedroom concentrations.

DISCUSSION

The traditional application of measurement error methods in regression makes use of a variety of assumptions concerning the independence of the structural and measurement components of variability. These assumptions are crucial in assuring the identifiability of the regression coefficients of interest. Frequently a physical basis can be found for the assumptions, as in the case of an instrument error due to a process unrelated to the latent variables being measured. In epidemiology, and many other applications, it is tempting to ascribe independence properties to many types of data inaccuracies without sufficient consideration of alternative models. In our example, the multiple indoor measurements in the Watertown data provide a means for evaluating the adequacy of a particular measurement error model. However, it would be far more satisfying if the Watertown data included at least a subset of individuals with both personal exposures and indoor measurements, thus permitting more direct tests of the conditional independence assumption.

The use of probit links leads to convenient closed forms for the induced model for the observed data. However, we have also examined logit links using approximate methods, with very similar results. This is not unexpected, due to the well known similarities between the two models, and we feel that the probit model, although not as frequently used with epidemiological data, has much to recommend it in measurement error and other latent variable models.

REFERENCES

1. Ferris, Jr., B. G., Speizer, F. E., Spengler, J. D., Dockery, D. W., Bishop, Y. M. M., Wolfson, J. M. and Humble, C. 'Effects of sulfur oxides and respirable particulates on human health: I. Methodology and demography of populations in study', American Review of Respiratory Disease, 120, 767-779 (1979).
2. Quackenboss, J. J., Spengler, J. D., Kanarek, S. M., Letz, R. and Duffy, C. P. 'Personal exposure to nitrogen dioxide: Relationship to indoor/outdoor air quality and activity patterns', Environmental Science & Technology, 20, 775-783 (1986).
3. Houthijs, D., Remijn, B., Brunekreef, B. and Koning, R. de 'Exposure to nitrogen dioxide and tobacco smoke and respiratory health of children', to appear in: Proceedings of the 4th International Conference on Indoor Air Quality and Climate, 17-21 August, Berlin(West), 1987.
4. Palmes, E. D., Gunnison, A. F., DiMattio, J. and Tomczyk, C. 'Personal sampler for NO₂', American Industrial Hygiene Association Journal, 37, 570-577 (1976).
5. McCullagh, P. 'Regression models for ordinal data' (with discussion), Journal of the Royal Statistical Society, Series B, 42, 109-142 (1980).
6. Burr-Doss, D. 'On errors-in-variables in binary regression-Berkson case', Technical Report No. 104, Division of Biostatistics, Stanford University, 1985.

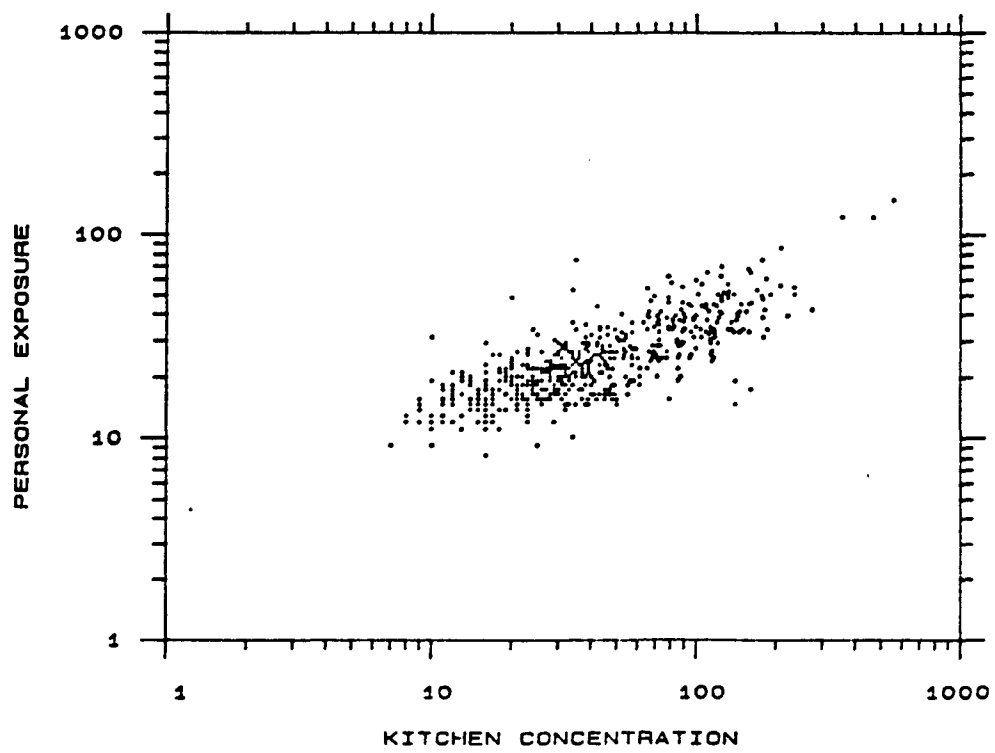
Table 1. Summary of regressions of log personal exposure on the log of indoor concentrations.

Data Set	Regression Coefficients						MSE
	Intercept		Kitchen		Bedroom		
	Est.	SE	Est.	SE	Est.	SE	
Dutch	1.17	0.06	0.30	0.02	0.33	0.03	0.05
Portage	1.28	0.10	0.28	0.07	0.33	0.07	0.07
Combined	1.22	0.05	0.30	0.02	0.33	0.02	0.06


FIGURE LEGENDS

Figure 1. Log personal exposure versus log kitchen concentration in the Dutch data set.

Figure 2. Log kitchen concentration versus log bedroom concentration in the health data set. Wheeze status is indicated by a 1 or a 0.



WATERTOWN DATA - 1 = WHEEZE, 0 = NO WHEEZE



>

Log kitchen NO2 (ug/l)

ACKNOWLEDGEMENTS

This work was supported by training grant HL07427 from the National Institutes of Health, research grants DMS-8702111 and DMS-8613681 from the National Science Foundation, and cooperative agreements CR-811650 and CR-813526 from the U. S. Environmental Protection Agency. We wish to thank Bert Brunekreef and Jim Quackenboss for their assistance in obtaining the data sets.