Ann. Occup. Hyg., Vol. 55, No. 1, pp. 97–112, 2011
Published by Oxford University Press
on behalf of the British Occupational Hygiene Society
doi:10.1093/annhyg/meq061

Analysis of Lognormally Distributed Exposure Data with Repeated Measures and Values below the Limit of Detection Using SAS

YAN JIN[1]*, MISTY J. HEIN[1], JAMES A. DEDDENS[1,2] and CYNTHIA J. HINES[1]

[1]Division of Surveillance, Hazard Evaluations and Field Studies, National Institute for Occupational Safety and Health, 4676 Columbia Pkwy, Cincinnati, OH 45226, USA; [2]Department of Mathematical Sciences, University of Cincinnati, Cincinnati, OH 45221, USA

Received 18 May 2010; in final form 13 July 2010; published online 20 December 2010

Studies of determinants of occupational exposure frequently involve left-censored lognormally distributed data, often with repeated measures. Left censoring occurs when observations are below the analytical limit of detection (LOD); repeated measures data results from taking multiple measurements on the same worker. A common method of dealing with this type of data has been to substitute a value (such as LOD/2) for the censored data followed by statistical analysis using the 'usual' methods. Recently, maximum likelihood estimation (MLE) methods have been employed to reduce bias associated with the substitution method. We compared substitution and MLE methods using simulated lognormally distributed exposure data subjected to varying amounts of censoring using two procedures available in SAS: LIFEREG and NLMIXED. In these simulations, the MLE method resulted in less bias and performed well even for censoring up to 80%, whereas the substitution method resulted in considerable bias. We illustrate the NLMIXED procedure using a dataset of chlorpyrifos air measurements collected from termiticide applicators on consecutive days over a 5-day workweek. We provide sample SAS code for several situations including one and two groups, with and without repeated measures, random slopes, and nested random effects.

Keywords: censoring; exposure determinants; limit of detection; LOD; NLMIXED; occupational exposure; repeated measures; SAS

INTRODUCTION

Statistical models are increasingly used to study the determinants of occupational exposure (e.g. Hines and Deddens, 2001; Lavoué *et al.*, 2005). In these models, the dependent variable is exposure, that is, the amount or concentration of a substance in a biological (e.g. exhaled breath, urine, or blood) or environmental (e.g. dermal or air) sample and the independent variables might include job characteristics, work practices, process information, personal protective equipment use, or engineering controls. Often the exposure data are right-skewed and can be assumed to follow a lognormal distribution (Leidel *et al.*, 1977). Statistical modeling of exposure data can be complicated when repeated measurements on workers are collected and in the presence of left-censored exposure data.

The importance of within- and between-worker variability in occupational exposure has long been recognized (Kromhout *et al.*, 1993) and the case for collecting repeat exposure measurements well documented (Rappaport *et al.*, 1995; Checkoway *et al.*, 2004). For example, when within-worker variability is high, multiple exposure measurements per worker are required to accurately characterize

*Author to whom correspondence should be addressed.
Tel: +1-513-841-4563; fax: +1-513-841-4486;
e-mail: hkn9@cdc.gov

a worker's exposure. Repeated measures data, however, can complicate modeling because measurements collected from the same worker or the same study site (when multiple study sites are studied) cannot be assumed to be independent. A second complication, left censoring, occurs when the measurement procedure has a limit of detection (LOD) and observations fall below this limit. The presence of data below the LOD has been termed Type I censoring (Lawless, 2003). For samples below the reported detection limit, laboratories typically provide no additional information. Furthermore, the data might be subject to multiple detection limits due to varying laboratory detection limits (i.e. detection limits may vary over time as analytical methods change, by batch or technicians within laboratories and across laboratories) or the introduction of multiple limits when laboratory results (e.g. milligrams of substance per sample) are expressed as a concentration (e.g. milligrams of substance per cubic meter of air collected). Left censoring can complicate analyses because traditional methods require a quantitative value for each observation. Statistical methods are readily available and widely used for analyzing repeated measures data without censoring (e.g. mixed-effects regression models) and for analyzing left-censored data without repeated measures (e.g. substitution methods, imputation methods, and tobit regression; Helsel, 1990); however, statistical methods for 'left-censored repeated measures data' (e.g. Pettitt, 1986; Hughes, 1999), while available, have not been widely adopted.

Regardless of whether the data include repeated measures, a convenient method of analyzing censored data has been to substitute a common value (e.g. LOD/2) for values below the LOD and proceed with the 'usual' statistical analysis. A review of methods used by occupational hygienists to evaluate determinants of exposure indicated wide use of the so called 'substitution method' (Burstyn and Teschke, 1999). Hornung and Reed (1990) showed, for nonrepeated measures data, that the LOD/2 substitution method should only be used when the data are highly skewed (i.e. GSD is $\sim \geq 3.0$) and recommended that substitution with LOD/$\sqrt{2}$ be used in other situations. Moreover, they used statistical simulation to show that substitution with LOD/2 or LOD/$\sqrt{2}$ could result in substantial bias when the proportion of censored data was large. For this reason, the use of maximum likelihood estimation (MLE) methods has been promoted (Helsel, 1990, 2006). MLE is a statistical method to fit parametric models and obtain estimates that maximize the likelihood of observing the collected data. The maximum likelihood (ML) estimator has various desirable properties, including consistency, asymptotic unbiasedness, and efficiency (van der Vaart, 2000). The mechanism of left censoring can be incorporated into parametric models; therefore, MLE can be used to estimate model parameters of left-censored data (Amemiya, 1973, 1984).

Historically, optimization of the likelihood function of censored data was often computationally intensive. For example, Hornung and Reed (1990) used the Hald method to calculate ML estimates of the mean and standard deviation of left-censored lognormal data. Rather than recommend its use, they presented the Hald method as a standard to compare the two substitution methods because the Hald method involved laborious calculations and the use of two different sets of tables. Advances in statistical computing have led to better algorithms for optimizing the censored likelihood function.

More recently, Hewett and Ganser (2007) reviewed various methods of estimating parameters (i.e. the mean and 95th percentile) of right-skewed left-censored exposure data. Based on statistical simulations, they concluded that the standard MLE method was the overall best method and mentioned that this method could now be easily implemented using the solver function of most computer spreadsheets (e.g. Finkelstein and Verma, 2001). Krishnamoorthy *et al.* (2009) proposed a model-based multiple imputation approach for analyzing sample data with nondetects. Flynn (2010) used the Microsoft Excel Solver tool to estimate the mean and standard deviation for several real exposure datasets subject to left censoring and Gillespie *et al.* (2010) described the use of reverse Kaplan–Meier estimators for estimating population means and percentiles from data with values below the LOD. All these methods, as currently presented, do not handle censored data within the repeated measures experimental design.

MLE using SAS procedures (SAS Institute Inc., Cary, NC, USA) has been used to analyze left-censored exposure data in a variety of settings. Slymen *et al.* (1994) illustrated the use of the LIFEREG procedure to fit simple regression models to left-censored normal data in an environmental bioaccumulation study that did not involve random effects. Thiébaut and Jacqmin-Gadda (2004) and Thiébaut *et al.* (2006) applied the NLMIXED procedure to longitudinal repeated measures lognormal data subject to left censoring in an infectious disease study. These methods, however, have not been widely adopted by occupational hygienists, possibly due to accessibility (Ogden, 2010). It is also not clear that how well they perform for large amounts of censoring.

In this paper, we outline the theory for left-censored lognormal data (with and without repeated measures) and summarize two SAS procedures that can be used to analyze left-censored exposure data using MLE methods, with emphasis on demonstrating the NLMIXED procedure to fit mixed-effects models to log-normally distributed occupational exposure data with repeated measurements. We present a simulation study to compare and evaluate the performance of estimates based on LOD/2 substitution and MLE methods for lognormally distributed data under a wide range of censoring levels. Finally, we illustrate the use of the NLMIXED procedure for the repeated measures model using a dataset of chlorpyrifos exposure measurements subject to left censoring at the LOD.

METHODS

Censored lognormal distribution

Let y be an observed exposure level following a lognormal distribution with geometric mean (GM) e^{μ} and geometric standard deviation (GSD) e^{σ}. The probability density function of y is given by

$$f(y|\mu,\sigma) = \frac{e^{\frac{-(\log y-\mu)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma y}, \quad y>0. \tag{1}$$

The probability of observing y less than the LOD (nondetect) is given by

$$P(y<\mathrm{LOD}) = \Phi\left(\frac{\log(\mathrm{LOD}) - \mu}{\sigma}\right), \tag{2}$$

where $\log(\cdot)$ denotes the natural logarithm and $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution with mean zero and standard deviation one. Therefore, the likelihood function for n independent observations, $y_{1,\ldots,}y_n$ can be written as

$$L(\mu,\sigma|y_1,\ldots,y_n) = \prod_1 f(y_{i_1}) \prod_0 P(y_{i_0}<\mathrm{LOD}), \tag{3}$$

where $\prod_1$ denotes the product over observed responses and $\prod_0$ denotes the product over censored responses. Amemiya (1973) proved that equation (3) is a valid likelihood function and that its maximizer holds all the properties of a ML estimator derived from complete data.

Censored repeated measures model

In many occupational exposure studies, repeated measures may be collected on an observational unit over time, space, or both time and space. For example, full-shift breathing zone air samples may be collected from study participants on consecutive work days, pre-, and post-shift biological samples may be collected from study participants each day over a workweek or area air samples may be collected at different locations on multiple days. The appropriate data model must recognize the relationship between repeated or serial observations on the same unit (Laird and Ware, 1982). Assuming that the underlying distribution of exposure measurements is lognormal, when the data includes both repeated measures and left censoring, a simple censored repeated measure model is given by

$$\log y^*_{ij} = X_{ij}\beta + u_i + \varepsilon_{ij}, u_i \sim N(0,\sigma^2_{\mathrm{b}}),$$
$$\varepsilon_{ij} \sim N(0,\sigma^2_{\mathrm{w}}), u_i \perp \varepsilon_{ij} \text{ for all } i \text{ and } j,$$

$$y_{ij} = \begin{cases} y^*_{ij} & \text{if } y^*_{ij} \geq \mathrm{LOD}_{ij}, \\ \text{missing} & \text{if } y^*_{ij} < \mathrm{LOD}_{ij}, \end{cases} \tag{4}$$

where y^*_{ij} is a latent response variable representing the true exposure level measured for subject i at time j with no censoring at the LOD; y_{ij} is the observed exposure level, detected at or above the LOD and censored below the LOD, $i = 1, 2, \ldots k$ (k is the number of subjects), $j = 1, 2, \ldots n_i$ (n_i is the number of observations for subject i); β is the vector of unknown population parameters (i.e. the vector of fixed effects); X_{ij} is the known design vector linking β to $\log y^*_{ij}$; u_i is the unknown individual effect (i.e. the random effect for subject i); ε_{ij} is the random error associated with subject i at time j; σ^2_{b} is the between-subject variance accounting for the variability of similar experimental units; and σ^2_{w} is the within-subject variance accounting for the variability of random errors in each measurement. In this model, u_i and ε_{ij} are assumed to be mutually independent and normally distributed and the resulting covariance structure compound symmetric.

The model specified in equation (4) is a conditional independence model, which means that given the random subject effects u_1, u_2, …, u_k, all responses y_{ij} are independent. Therefore, based on equation (3) and conditional on the subject effects, the likelihood function for the censored repeated measures model is given by where

$$L(\beta,\sigma_{\mathrm{w}},\sigma_{\mathrm{b}}|u_i,y_{ij},i=1,\ldots,k,j=1,\ldots,n_i)$$
$$= \prod_i \prod_j g(y_{ij}|\beta,\sigma_{\mathrm{w}},u_i), \tag{5}$$

$$g(y_{ij}|\beta,\sigma_w,u_i)=\begin{cases}\dfrac{e^{\frac{-(\log y_{ij}-X_{ij}\beta-u_i)^2}{2\sigma_w^2}}}{\sqrt{2\pi}\sigma_w y_{ij}} & \text{if } y_{ij} \text{ is observed,}\\ \Phi\left(\dfrac{\log(\mathrm{LOD}_{ij})-X_{ij}\beta-u_i}{\sigma_w}\right) & \text{if } y_{ij} \text{ is censored.}\end{cases}$$

Since the probability density of u_i is fully determined by σ_b, i.e. $q(u_i|\sigma_b)=\dfrac{e^{\frac{-u_i}{2\sigma_b^2}}}{\sqrt{2\pi}\sigma_b}$, it is possible to derive the marginal likelihood function (McCullagh and Nelder, 1989) for the model specified in equation (4) by integrating out u_i from equation (5), resulting in

$$m(\beta,\sigma_w,\sigma_b|y_{ij},i=1,\ldots,k,j=1,\ldots,n_i)=\prod_i\int\prod_j g(y_{ij}|\beta,\sigma_w,u_i)q(u_i|\sigma_b)\mathrm{d}u_i, \quad (6)$$

Finally, the ML estimates of β, σ_w, σ_b can be obtained by minimizing the negative of the logarithm of equation (6).

Use of SAS procedures in the analysis of censored data

The SAS procedure LIFEREG can be used to perform MLE for left-censored data with no repeated measures. Two new variables, called lower and upper endpoints, must be created for each observation in the dataset. For outcomes that were detected (i.e. not censored), the lower and upper endpoints are set to the observed levels. For censored outcomes (below the LOD), the lower endpoint is set to 'missing' and the upper endpoint is set to the value of the LOD. In the context of lognormally distributed data, the missing value represents negative infinity, the theoretical lower bound for the log-transformed data. For example, for the one sample model, consider a SAS dataset ONE with variable *Y* representing the exposure level, *C* indicating whether the observation was censored, and LOD representing the LOD associated with each observation. To estimate the GM and GSD, create a new dataset TWO with variables, YL and YU, as illustrated below:

```
DATA TWO;
  SET ONE;
  *comment: here C=1 indicates that Y is below
   the LOD;
  IF C=1 THEN DO; YL=.; YU=LOD; END;
  IF C=0 THEN DO; YL=Y; YU=Y; END;
  RUN;
```

The data are now in the appropriate format to perform MLE. Assuming a lognormal distribution and that no covariates are included in the model, the following SAS code will generate the ML estimates of the GM and GSD:

```
PROC LIFEREG DATA=TWO;
  MODEL (YL,YU)=/D=LOGNORMAL;
  ODS OUTPUT PARAMETERESTIMATES=
  PEY;
  RUN;
DATA PEY2;
  SET PEY;
  IF PARAMETER='Intercept' THEN GM=
  EXP(ESTIMATE);
  IF PARAMETER='Scale' THEN GSD=
  EXP(ESTIMATE);
  RUN;
PROC PRINT DATA=PEY2;
  VAR GM GSD;
  RUN;
```

For the two-group model, if the dataset included an indicator variable for group (i.e. $X = 0$ for Group 1, $X = 1$ for Group 2), changing the model statement to the following would generate group-specific GMs and an overall GSD:

```
MODEL (YL,YU) = X/D = LOGNORMAL;
```

In this situation, the Group 1 GM would be obtained by exponentiation of the intercept and the Group 2 GM would be obtained by exponentiation of the sum of the intercept and the Group 2 estimate. The overall GSD would be obtained by exponentiation of the scale parameter. For either model, additional covariates could be added to the MODEL statement, if desired.

The SAS procedure NLMIXED can be used to perform MLE for repeated measures data subject to left censoring (Thiébaut and Jacqmin-Gadda, 2004). NLMIXED requires specification of the conditional likelihood function of the censored data, given the random effects, as in equation (5). For example, for SAS dataset THREE, assume that variables *Y*, LOD, and *C* are as before but that the data includes repeated measurements on each subject, identified by the variable SUBJECTID (i.e. the random effect). The SAS program to fit the one-group censored repeated measures model to lognormally distributed data is given by the following:

```
DATA FOUR;
  SET THREE;
  *comment: here C=1 indicates Y is below the
  LOD;
  IF C=1 THEN NY=LOD;
```

```
  IF C=0 THEN NY=Y;
  RUN;
PROC NLMIXED DATA=FOUR;
  PARMS BETA0=0 SB_2=2 SW_2=2;
  BOUNDS SB_2 SW_2>0;
  MU=BETA0 + U_I;
  PI=2*ARSIN(1);
  IF C=0 THEN L=(1/(SQRT(2*PI*SW_2)*NY))
  *EXP(-(LOG(NY)-MU)**2/(2*SW_2));
  IF C=1 THEN L=PROBNORM((LOG(NY)-MU)/
  SQRT(SW_2));
  LL=LOG(L);
  MODEL NY ~ GENERAL(LL);
  RANDOM U_I ~ NORMAL(0,SB_2) SUBJECT=
  SUBJECTID;
  RUN;
```

In the NLMIXED procedure, the conditional distribution is constructed using SAS programming statements; the GENERAL function used in the MODEL statement specifies a general log-likelihood function that is defined by the user instead of following a standard distribution form; the optional PARMS and BOUNDS statements can be used to specify initial values and boundary constraints for the parameters. Including NY in the denominator of the likelihood function for noncensored observations is not necessary because it only contributes a constant term to the log-likelihood (note that Thiébaut and Jacqmin-Gadda (2004) did not include NY in their program). The estimated coefficients will be the same regardless of whether NY is included; however, doing so gives the correct likelihood. Here, BETA0 is the estimate of the log of the GM and exp[sqrt(SB_2 + SW_2)] is the estimate of the total GSD. Furthermore, the NLMIXED procedure can be used in the absence of repeated measures by omitting the RANDOM statement and excluding the random effect U_I in the definition of MU, in which case NLMIXED and LIFEREG give the same estimates and standard errors.

For the two-group censored repeated measures model, if variable X is an indicator variable for Group 2, only the PARMS statement and the definition of MU need to be changed:

```
PARMS BETA0=0 BETA1=0 SB_2=2 SW_2=2;
MU=BETA0 + BETA1*X + U_I;
```

As before, additional covariates could be included in the definition of MU to adjust for confounders.

Simulation study

The goal of the simulation study was to evaluate the performance of the MLE method and to compare it with the common substitution method in analysis of lognormally distributed exposure data with values below the LOD. All simulations and analyses were done using SAS version 9.2.

In the first simulation, we generated data from a two-group repeated measures design with three repeated observations per subject and no time effect (i.e. a compound symmetric covariance structure). The censored repeated measures model was given by

$$\log y_{ij}^{*} = \beta_0 + \beta_1(\text{group}_i = 2) + u_i + \varepsilon_{ij},$$

$$y_{ij} = \begin{cases} y_{ij}^{*} & \text{if } y_{ij}^{*} \geq \text{LOD}_{ij}, \\ \text{missing} & \text{if } y_{ij}^{*} < \text{LOD}_{ij}, \end{cases} \qquad (7)$$

where y_{ij} is the jth measurement for the ith subject, $u_i \sim N\left(0, \sigma_{\text{b}}^2\right)$, and $\varepsilon_{ij} \sim N\left(0, \sigma_{\text{w}}^2\right)$. GMs for Groups 1 and 2 were given by $\exp(\beta_0)$ and $\exp(\beta_0 + \beta_1)$, respectively. Thus, β_1 was the logarithm of the ratio of the group GMs. The total GSD was given by $\exp\left[\text{sqrt}\left(\sigma_{\text{b}}^2 + \sigma_{\text{w}}^2\right)\right]$. We evaluated three different sample sizes (15, 30, and 100 subjects per group), two total GSDs (three and five), and three ratios of the between- to within-subject variance on the log scale (1:5, 1:1, and 5:1). We specified a Group 1 GM of 200 [i.e. $\beta_0 = \ln(200)$] and a Group 2 GM of 400 [i.e. $\beta_1 = \ln(400/200) = \ln(2)$]. These GMs were selected based on reasonable power ($\sim$0.80) of detecting a difference between the groups in the scenario with 30 subjects per group, a GSD of 3, and equal between- and within-subject variances. Each dataset was subjected to censoring (0–80% censoring) and analyzed using two methods: the MIXED procedure with LOD/2 substitution for any censored observations [using the default restricted ML (REML) estimation method] and the NLMIXED procedure based on the MLE method (as described above). Summary statistics from 1000 datasets for each of the $3 \times 2 \times 3 = 18$ scenarios included mean percent bias for β_0, β_1, σ_{b}^2, and σ_{w}^2 (mean percent difference between the estimate and the true parameter value), the empirical 5% Type I error rate for β_1 [the fraction of simulations for which the 95% confidence interval (CI) did not contain the true value of β_1] and the empirical power for β_1 (the fraction of simulations for which the 95% CI did not contain the null value of 0).

A second simulation was conducted to evaluate the MLE method in fitting simple regression models to censored lognormally distributed data in the absence of repeated measures. As before, we generated data for two groups and evaluated three different sample sizes (15, 30, and 100 subjects per group) and two GSDs (three and five). Here, the GSD was given by exp[sqrt(σ^2)]. We specified a Group 1

GM of 200 [i.e. $\beta_0 = \ln(200)$] and a Group 2 GM of 500 [i.e. $\beta_1 = \ln(500/200) = \ln(2.5)$]. These GMs were selected based on reasonable power (∼0.80) of detecting a difference between the groups in the scenario with 30 subjects per group and a GSD of 3. Each dataset was subjected to censoring (0–80% censoring) and analyzed using two different methods: the GLM procedure with LOD/2 substitution for any censored observations and the LIFEREG procedure based on the MLE method (as described above). Summary statistics from 1000 datasets for each of the $3 \times 2 = 6$ scenarios included the mean percent bias for β_0, β_1, and σ^2, the empirical 5% Type I error rate for β_1, and the empirical power for β_1.

RESULTS

Simulation results evaluating the ML and substitution methods with respect to parameter estimation (on the log scale) in the two-group repeated measures model for lognormally distributed data with 0–80% censoring are summarized in Fig. 1. For β_0 (the logarithm of the Group 1 GM), little bias was observed for the ML estimates but the LOD/2 substitution method overestimated the Group 1 GM for censoring rates >25% (GSD = 3) or 10% (GSD = 5). The amount of bias for the LOD/2 substitution method was independent of the sample size and ratio of between- to within-subject variance but not the total variance. For β_1 (the difference of the logarithms of the group GMs), the LOD/2 substitution method produced markedly downward bias when the percentage of censoring was high, whereas the ML estimates were quite stable even for censoring levels up to 80%. For zero censoring, the ML estimates of the between-subject variance were biased, particularly for small sample sizes, whereas the LOD/2 substitution method (REML estimation) was not. However, in the presence of censoring, the ML estimates of the between-subject variance were less biased compared to the LOD/2 substitution method (REML estimation). ML estimates of the within-subject variance were unbiased, regardless of censoring; whereas, the LOD/2 substitution method was biased. For all parameters, the amount of bias associated with the LOD/2 substitution method was higher when the GSD was larger.

The empirical 5% Type I error rates for hypotheses about β_1 (the difference in the logs of the GMs) was close to 5% for both methods as long as censoring rates did not exceed 25% (Fig. 2). Above 25% censoring, the error rate of the LOD/2 substitution method rapidly increased; whereas the error rate of the MLE method remained at ∼5%. Note that larger sample sizes produced higher error rates for the LOD/2 method due to narrower CIs. In contrast to the bias in estimating β_1 and the empirical Type I error rate for β_1, the power of the test for β_1 was similar for the two methods. This suggests that if one is more interested in detecting group effect differences rather than estimating group means, the substitution method performs as well as the MLE method for censoring levels up to 80% for the sample sizes, GSDs, and ratios of between- and within-subject variances evaluated. However, for both methods, a loss of power was observed for high levels of censoring, particularly for small sample sizes. As expected, higher power was associated with larger sample sizes and smaller GSDs. Higher power was also associated with a smaller ratio of between- to within-subject variances as the variance of the group effect is a monotone increasing function of the between-subject variance. $(\mathrm{var}(\hat{\beta}_1) = \mathrm{var}\left(\frac{1}{45}\sum_{i=1}^{15}\sum_{j=1}^{3}\log Y_{ij} - \frac{1}{45}\sum_{i=16}^{30}\sum_{j=1}^{3}\log Y_{ij}\right)$ $= 2\frac{\sigma_b^2}{15} + 2\frac{\sigma_w^2}{45} = \frac{2}{45}((\sigma_b^2 + \sigma_w^2) + 2\sigma_b^2)$, for 15 subjects per group and three repeats per subject). Similarly, the rejection rate was inflated in the LOD/2 method due to tighter confidence bands produced by the smaller ratio of variance components.

Similar results were observed in the simple two-group models in the absence of repeated measures (Figs. 3 and 4). Biases in estimating the group means, the mean difference, and the variance were greater based on the LOD/2 substitution method compared to the MLE method using the LIFEREG procedure. As expected, when the sample size was small, we observed bias in the estimation of variance based on the MLE method.

EXAMPLE

In 1998, the National Institute for Occupational Safety and Health conducted a study of termite control workers who applied chlorpyrifos-containing termiticides to residential and commercial structures in North Carolina from early March to early July (Hines and Deddens, 2001). Thirty-seven applicators, all males, participated in the study. Full-shift breathing zone air samples were collected from each applicator on consecutive days during a 5-day workweek. Information regarding treatment of enclosed crawl spaces was noted for each day. A total of 184 air samples (four to five for each applicator) were collected for determination of chlorpyrifos

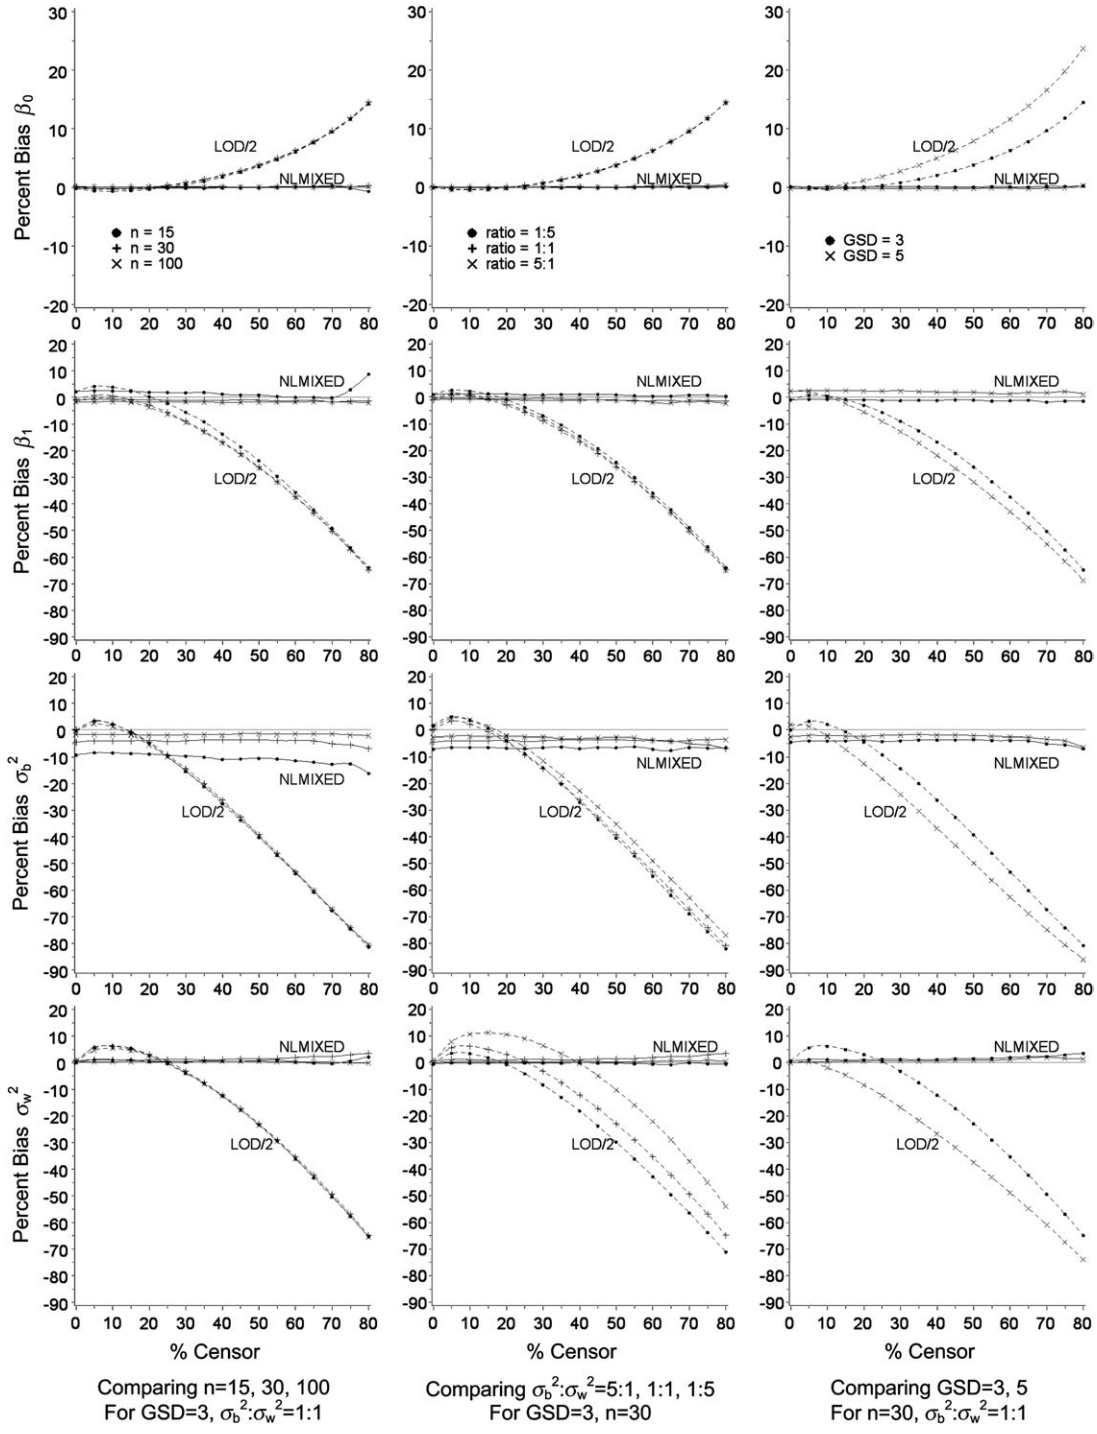

Fig. 1. Simulation results for the two-group repeated measures design comparing analyses using LOD/2 substitution (MIXED procedure with REML) and MLE (NLMIXED procedure) methods for lognormally distributed data. β_0 is the logarithm of the Group 1 GM [i.e. $\beta_0 = ln(200)$]; β_1 is the difference in the logarithms of the group GMs [i.e. $\beta_1 = ln(GM_2/GM_1) = ln(2)$]; σ_b^2 is the between-worker variance of the log-transformed values; σ_w^2 is the within-worker variance of the log-transformed values; and GSD = exp[sqrt($\sigma_b^2+\sigma_w^2$)]. Percent bias is the mean percent difference between the modeled estimate and the true parameter value for 1000 simulated datasets.

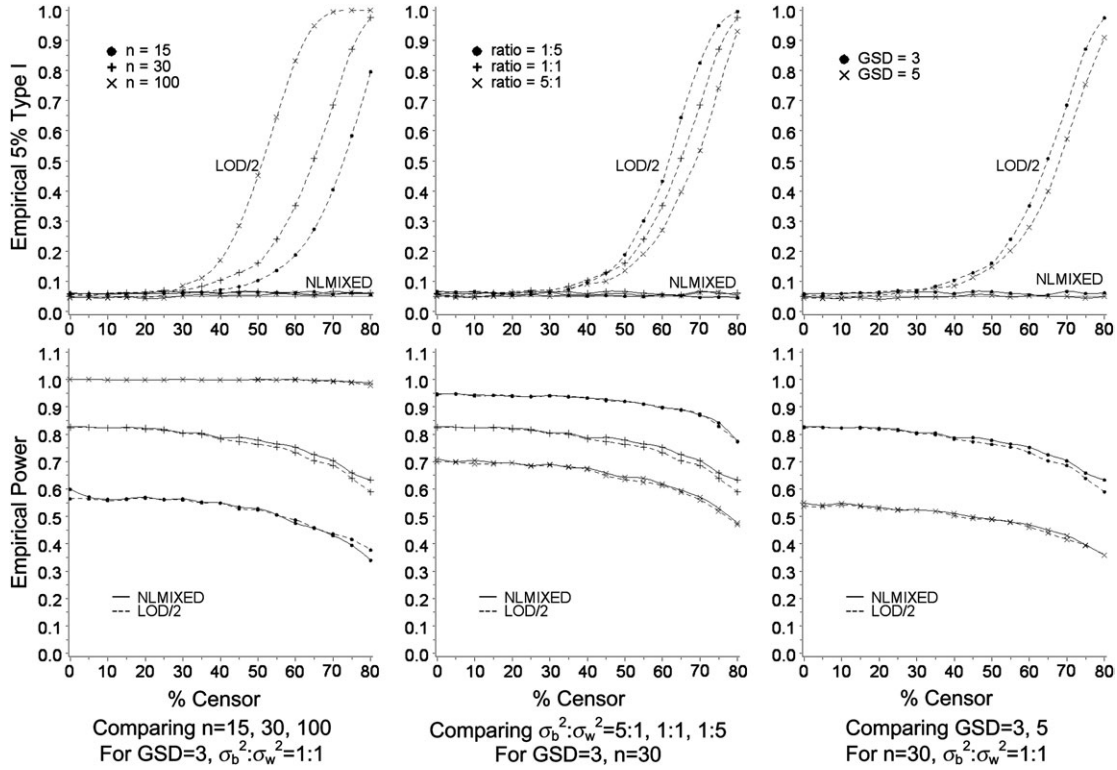

Fig. 2. Simulation results for the two-group repeated measures design comparing analyses using LOD/2 substitution (MIXED procedure with REML) and MLE (NLMIXED procedure) methods for lognormal data. β_1 is the difference in the logarithms of the group GMs [i.e. $\beta_1 = ln(GM_2/GM_1) = ln(2)$]; σ_b^2 is the between-worker variance of the log-transformed values; σ_w^2 is the within-worker variance of the log-transformed values; and GSD = exp[sqrt($\sigma_b^2 + \sigma_w^2$)]. The empirical 5% Type I error rate is the fraction of the 1000 simulated datasets for which the 95% CI for β_1 did not contain the true value of β_1. The empirical power is the fraction of the 1000 simulated datasets for which the 95% CI for β_1 did not contain the value of 0.

exposure levels. Chlorpyrifos detection limits ranged from 0.05 to 0.20 μg per sample; reported chlorpyrifos levels ranged from nondetect to 73 μg per sample. Laboratory data (micrograms per sample) were converted to concentrations in air (micrograms per cubic meter) by dividing by the sample volumes. The distribution of the resulting concentrations was highly right-skewed and a visual examination of the QQ plot indicated that the data were consistent with lognormality; therefore, a natural logarithm transformation was applied to the chlorpyrifos air concentrations prior to statistical analyses.

The resulting data are presented in Table 1 for four selected workers. The original data were subject to very little censoring with only 3 of 184 samples censored at the mass-based LOD. For example, all of worker A's measurements were above the laboratory LOD, but two of worker C's measurements were reported as nondetect. To evaluate methods of analyzing censored repeated measures data, the chlorpyrifos mass data were subjected to censoring at three levels: 20, 40, and 60% based on the 20th, 40th, and 60th percentiles of the chlorpyrifos mass, respectively. Under hypothetical censoring at 40%, two of worker A's measurements (4.5 and 3.0 μg) and all of worker C's measurements were <4.6 μg, the 40th percentile of the reported chlorpyrifos mass. Since the air volume varied from sample to sample, the censored concentrations ranged from <7.0 to <110 μg m^{-3} for these four workers.

We considered a two-group repeated measures model by including an indicator variable for crawl space treatment. The original and censored datasets were each analyzed using two methods. Method 1 replaced nondetects with one-half the censoring level (i.e. one-half the mass-based LOD divided by the air volume) and used the MIXED procedure (REML) to estimate the model parameters; Method 2 censored nondetects at the censoring level and used the NLMIXED procedure (ML) to estimate

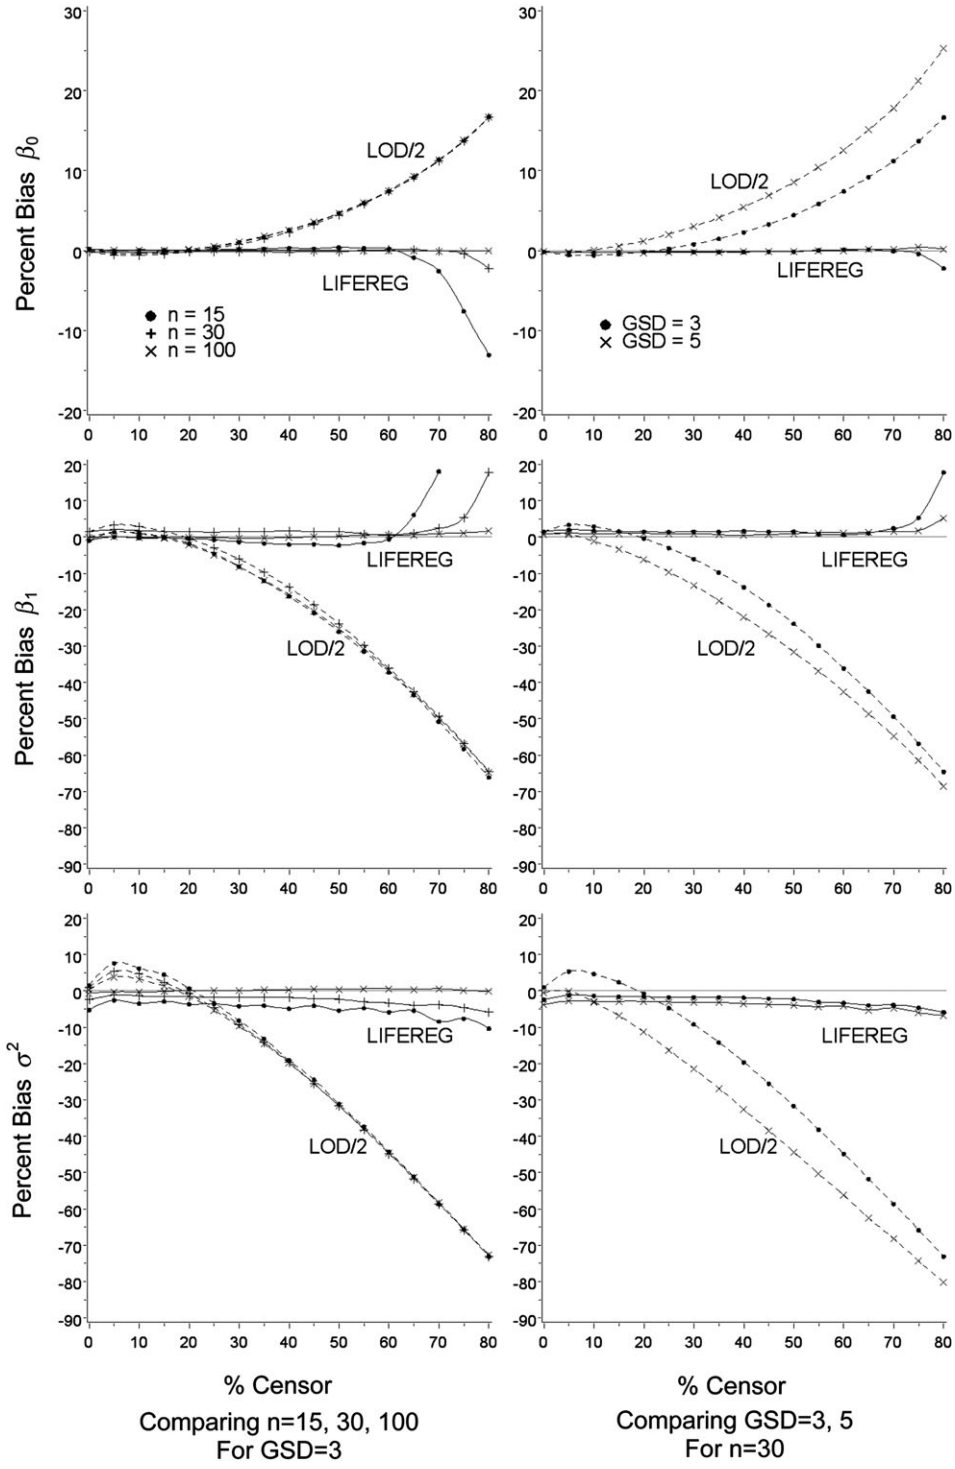

Fig. 3. Simulation results for the simple two-group model (no repeated measures) comparing analyses using LOD/2 substitution (GLM procedure) and MLE (LIFEREG procedure) methods for lognormal data. β_0 is the logarithm of the group 1 GM [i.e. $\beta_0 = ln(200)$]; β_1 is the difference in the logarithms of the group GMs [i.e. $\beta_1 = ln(GM_2/GM_1) = ln(2.5)$]; σ^2 is the log transformed variance; and GSD=exp[sqrt(σ^2)]. Percent bias is the mean percent difference between the modeled estimate and the true parameter value for 1000 simulated datasets.

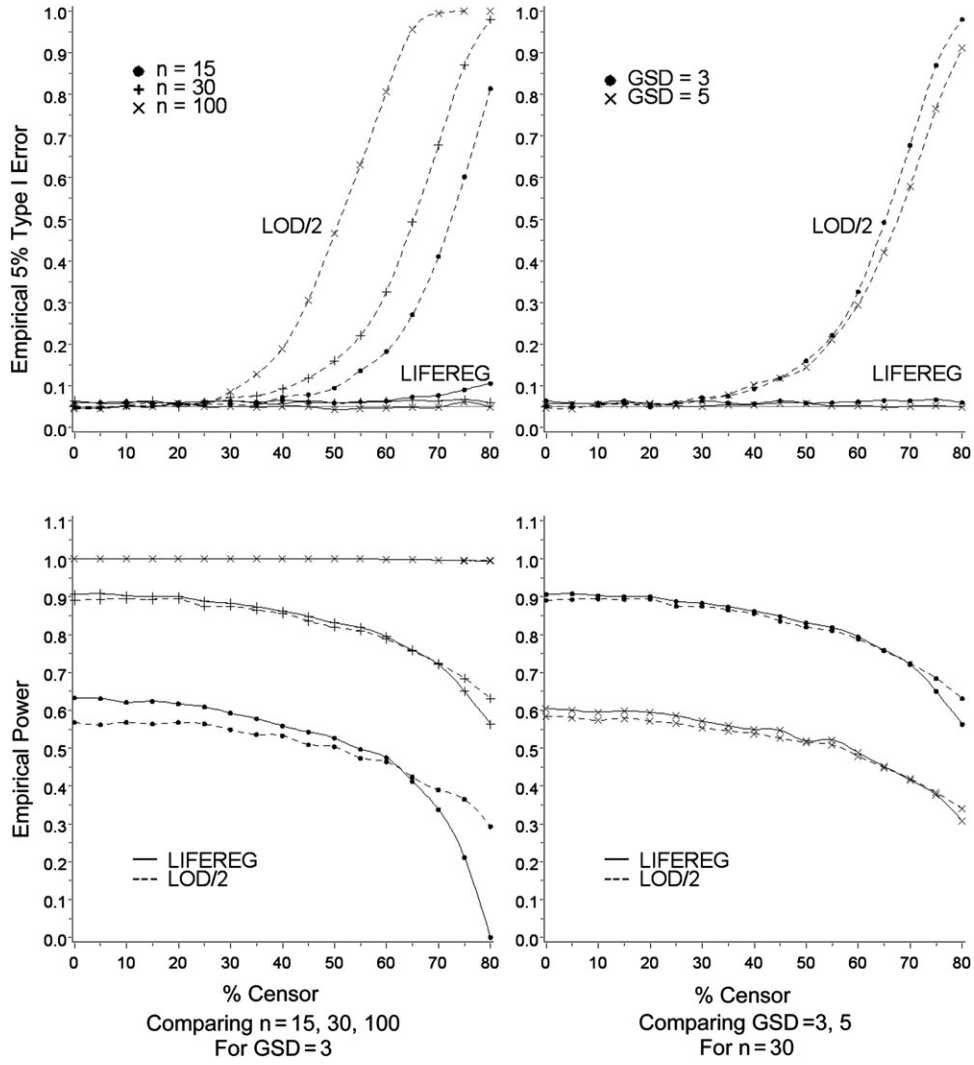

Fig. 4. Simulation results for the simple two-group model (no repeated measures) comparing analyses using LOD/2 substitution (GLM procedure) and MLE (LIFEREG procedure) methods for lognormal data. β_1 is the difference in the logarithms of the group GMs [i.e. $\beta_1 = ln(GM_2/GM_1) = ln(2.5)$]. The empirical 5% Type I error rate is the fraction of the 1000 simulated datasets for which the 95% CI for β_1 did not contain the true value of β_1. The empirical power is the fraction of the 1000 simulated datasets for which the 95% CI for β_1 did not contain the value of 0.

the model parameters. SAS code is provided in Program A. The resulting parameter estimates are provided in Table 2. Since the original dataset was largely uncensored, Methods 1 and 2 produced similar estimates. Under hypothetical censoring, Method 1 (LOD/2 substitution) produced 95% CIs that did not contain the uncensored estimates for β_0, β_1, and σ_b^2 at 40% censoring and for all four parameters at 60% censoring; whereas 95% CIs based on Method 2 (NLMIXED) contained the uncensored estimates for all parameters at all censoring levels evaluated.

DISCUSSION

We have illustrated the use of the NLMIXED procedure in SAS to perform repeated measures analysis of variance when the dependent variable is lognormally distributed and left censored due to values below the LOD. Clearly, the MLE method outperformed the substitution method. Of all the factors in the simulations, including percent censoring, sample size, GSD, and the ratio of variance components, sample size was the most important factor influencing the performance of the MLE method,

Table 1. Original laboratory chlorpyrifos air sample data and hypothetical data under mass-based censoring for 4 of 37 selected termiticide applicators sampled daily over a 5-day workweek[a]

Day	Crawl space indicator	Air volume (l)	Detection limit (μg)	Original laboratory data[b]		Hypothetical data under mass-based censoring[c]					
						20% at 1.7 μg		40% at 4.6 μg		60% at 8.6 μg	
				Mass (μg)	Concentration (μg m^{-3})	Mass (μg)	Concentration (μg m^{-3})	Mass (μg)	Concentration (μg m^{-3})	Mass (μg)	Concentration (μg m^{-3})
Subject A											
M	1	549	0.10	4.5	8.2	4.5	8.2	**<4.6**	**<8.4**	**<8.6**	**<16**
Tu	0	589	0.10	3.0	5.1	3.0	5.1	**<4.6**	**<7.8**	**<8.6**	**<15**
W	0	568	0.10	9.5	17	9.5	17	9.5	17	9.5	17
Th	0	570	0.10	23	41	23	41	23	41	23	41
F	1	560	0.10	37	66	37	66	37	66	37	66
Subject B											
M	1	624	0.07	4.5	7.2	4.5	7.2	**<4.6**	**<7.4**	**<8.6**	**<14**
Tu	1	670	0.07	9.2	14	9.2	14	9.2	14	9.2	14
W	0	616	0.07	3.2	5.2	3.2	5.2	**<4.6**	**<7.5**	**<8.6**	**<14**
Th	1	635	0.07	9.0	14	9.0	14	9.0	14	9.0	14
F	0	297	0.07	0.43	1.4	**<1.7**	**<5.7**	**<4.6**	**<16**	**<8.6**	**<29**
Subject C											
M	0	473	0.05	0.28	0.6	**<1.7**	**<3.6**	**<4.6**	**<9.7**	**<8.6**	**<18**
Tu	0	41	0.05	**<0.05**	**<1.2**	**<1.7**	**<41**	**<4.6**	**<110**	**<8.6**	**<210**
W	0	511	0.05	1.7	3.3	**<1.7**	**<3.3**	**<4.6**	**<9.0**	**<8.6**	**<17**
Th	0	655	0.05	1.6	2.4	**<1.7**	**<2.6**	**<4.6**	**<7.0**	**<8.6**	**<13**
F	0	50	0.05	**<0.05**	**<1.0**	**<1.7**	**<34**	**<4.6**	**<92**	**<8.6**	**<170**
Subject D											
M	1	524	0.20	21	40	21	40	21	40	21	40
Tu	1	628	0.20	15	24	15	24	15	24	15	24
W	0	803	0.20	11	14	11	14	11	14	11	14
Th	1	769	0.20	67	88	67	88	67	88	67	88
F	1	428	0.20	16	37	16	37	16	37	16	37

M, Monday; Tu, Tuesday; W, Wednesday; Th, Thursday; F, Friday.
[a]From Hines and Deddens (2001).
[b]Original laboratory data reported as chlorpyrifos mass (microgram) was converted to a chlorpyrifos concentration by dividing by the sample air volume (cubic meter). Samples reported as 'nondetect' are displayed in bold as '$<LOD_{mass}$' and '$<LOD_{concentration}$'.
[c]Hypothetical data under mass-based censoring considered the 20th, 40th, and 60th percentiles of the reported chlorpyrifos mass: 1.7, 4.6, and 8.6 μg, respectively. Samples with mass values less than these are displayed in bold as '$<LOD_{mass}$' and '$<LOD_{concentration}$'.

especially for estimating the between-subject variance. Bias in estimating the between-subject variance was highest for smaller sample sizes. Furthermore, ML estimates of the between-subject variance were biased even in the absence of censoring, especially for small sample sizes, whereas REML estimates were not biased. This results from the fact that ML estimators are only asymptotically unbiased and because REML estimation takes into account the loss in degrees of freedom associated with the group effect when estimating the between-subject variance (Swallow and Monahan, 1984). A recent commentary in the *Annals of Occupational Hygiene* strongly opposed the use of substitution methods and urged journals to 'consider it a flawed method compared to other methods that are available and to reject papers that use it' (Helsel, 2010). Although we support the use of nonsubstitution methods, we believe there may be some exceptions. For example, when the primary interest is estimating the between-subject variance for small sample sizes with low censoring (<15%), REML estimation (and

Table 2. Results of analysis of censored repeated measures exposure data from 37 termiticide applicators[a]

Parameter[b]	Percent censoring	Method 1: LOD/2 substitution MIXED–REML		Method 2: MLE NLMIXED	
		Estimate	95% CI	Estimate	95% CI
β_0	1.6	1.95	1.62–2.29	1.96	1.63–2.29
	20	2.06	1.75–2.37	2.00	1.67–2.34
	40	2.30	2.04–2.55[c]	2.15	1.82–2.47
	60	2.56	2.35–2.76[c]	2.22	1.84–2.59
β_1	1.6	0.77	0.52–1.03	0.77	0.51–1.03
	20	0.68	0.43–0.93	0.72	0.47–0.97
	40	0.46	0.23–0.69[c]	0.60	0.33–0.88
	60	0.29	0.08–0.49[c]	0.55	0.25–0.85
σ_b^2	1.6	0.77	0.48–1.43	0.73	0.46–1.36
	20	0.60	0.37–1.15	0.75	0.46-1.46
	40	0.35	0.21–0.72[c]	0.55	0.32-1.18
	60	0.22	0.12–0.46[c]	0.56	0.30–1.43
σ_w^2	1.6	0.62	0.50–0.79	0.59	0.48–0.76
	20	0.58	0.47–0.75	0.51	0.40–0.67
	40	0.52	0.42–0.67	0.49	0.37–0.68
	60	0.40	0.32–0.51[c]	0.44	0.31–0.66

[a]Data from Hines and Deddens (2001).
[b]The repeated measures model was given by $y_{ij}=\beta_0+\beta_1(X_{ij}=1)+u_i+\epsilon_{ij}$, where y_{ij} is the natural log-transformed chlorpyrifos concentration and X_{ij} is an indicator for crawl space for subject i and sample j, u_i is $N(0,\sigma_b^2)$, ϵ_{ij}, is $N(0,\sigma_w^2)$, and u_i and ϵ_{ij} are independent.
[c]95% CI does not contain the 'uncensored' estimate. Note that 95% CIs for variance components were constructed using the Satterthwaite approximation with a lower boundary constraint of zero (Milliken and Johnson, 1992).

the LOD/2 substitution) performed better in our simulations than the MLE method. For all levels of censoring, however, the MLE method performed better than LOD/2 substitution with REML for estimating the group means, the difference in group means, and within-subject variances. MLE also performed better than LOD/2 substitution with REML for estimating the between-subject variances for censoring $>\sim$25%.

We recognize that for some more complicated models, software and statistical support for analyzing left-censored repeated measures data may not yet be available. For example, we know of no methods for analyzing left-censored repeated measures data when the appropriate covariance structure is first-order autoregressive. The NLMIXED procedure at this time does not accommodate covariance structures more complicated than compound symmetry because the special covariance matrix structure of the random error term at the subject level cannot be specified. We would encourage analysts that are tempted to use the LOD/2 substitution method (e.g. because no other method is available or because the amount of censoring is very low) to first evaluate its performance using simulation methods such as those presented here.

In our simulations, the two groups shared a common within-worker variance, but different within-worker variances can also be specified using the NLMIXED procedure, as illustrated in Program B. Also, in our simulations, the covariate group was binary. For a continuous covariate, Thiébaut and Jacqmin-Gadda (2004) showed how to specify a model with both a random intercept and a random slope in the presence of censoring. Sample SAS code for this situation is provided in Program C. Note that unlike the MIXED procedure, the NLMIXED procedure does not have a CLASS statement, so categorical independent variables must be converted to binary indicator variables in a prior DATA step. The NLMIXED procedure also does not have a method for adjusting for multiple comparisons; however, a Bonferroni adjustment could be applied to *P* values after estimating the effect differences among categorical levels.

The NLMIXED procedure can also allow more than one level of random effects to be included in the repeated measures model (Littell *et al.*, 2006). SAS code to fit nested random effects models for simulated normally distributed data are provided in Program D; however, the reader is cautioned that even for small or moderately sized datasets, in the presence

Table 3. Central processing unit time (h = hour, m = minute, and s = second) required by the NLMIXED procedure to fit nested random effects models for simulated normal data ($\sigma_a = 3$, $\sigma_b = 2$, and $\sigma_w = 1$) with 30% nondetects. Programs were executed on a personal computer equipped with an Intel Core 2 Duo processor at 3.0 GHz and 2 GB of memory. Results were based on a single dataset simulated under each combination of number of levels of random factors A and B nested within A, with three repeated measurements per subject

Number of levels of Factor A	Number of levels of Factor B (nested within A)			
	2	3	4	5
10	5s	1 m 00 s	12 m 19 s	3 h 38 m 17 s
20	9s	2 m 12 s	25 m 35 s	5 h 54 m 08 s
40	14s	8 m 35 s	42 m 56 s	8 h 17 m 57 s

of censoring and nested random effects NLMIXED may take a long time to run. We compared run times for a nested random effects design with three repeats per subject in the presence of 30% censoring on a personal computer equipped with an Intel Core 2 Duo processor at 3.0 GHz and 2 GB of memory. Central processing unit time varied with the number of levels of the first random effect (A) and (mainly with) the number of levels of the second random effect (B nested within A) (Table 3). Another limitation is that NLMIXED requires the number of levels of the first random effect (A) be strictly greater than the number of levels of the second random effect (B) plus one. This is due to the fact that NLMIXED uses approximate degrees of freedom for *t*-values associated with parameter estimates, computed as the number of subjects minus the number of random effects specified in the random statement. Consequently, a study with a random effect for study site, a nested random effect for worker, and with repeated measurements and censoring could only be analyzed using NLMIXED if the number of workers sampled from each study site was less than the number of sites minus one. If the number of workers at each study site exceeded the number of sites, NLMIXED would generate invalid parameter estimates with negative degrees of freedom.

Our study has some limitations. First, our simulations were performed assuming a lognormal distribution and we did not evaluate the robustness of these methods to departures from lognormality. In practice, researchers should always employ graphical or other known methods for testing the lognormal assumption. The NLMIXED procedure, however, allows the user to explicitly specify the likelihood function, so other censored distributions, including the censored truncated normal distribution, can be evaluated. Second, in our simulations, we assumed that a single LOD applied to all measurements. Both LIFEREG and NLMIXED allow the LOD to vary by measurement. We expect similar conclusions for multiple LODs but did not specifically evaluate this. Last, we must emphasize that this article presents a method for handling repeated measures analysis of variance data with values below the LOD that uses MLE, consequently it is only asymptotically unbiased.

CONCLUSIONS

SAS software can be easily used to analyze left-censored lognormal exposure data using MLE methods for simple models (LIFEREG procedure) or repeated measures models (NLMIXED procedure). Overall, MLE is a better method than LOD/2 substitution for parameter estimation and hypothesis testing. In fact, these simulations have empirically shown that the MLE method provides approximately unbiased estimates for the sample sizes considered and even for censoring rates >50%. We recommend the use of MLE methods for the analysis of left-censored exposure data with one exception (when the primary interest is estimation of the between-subject variance for small sample sizes with low censoring (<15%), in which case REML estimation (and LOD/2 substitution) may be used.

FUNDING

This work was supported by intramural research funds from the National Institute for Occupational Safety and Health.

Disclaimer—The findings and conclusions in this report are those of the authors and do not necessarily represent the views of the National Institute for Occupational Safety and Health.

APPENDIX

PROGRAMS

All SAS code are available in a supplemental text file that can be copied and pasted directly into SAS (see *Annals of Occupational Hygiene* online for details).

Program A:

SAS code for the chlorpyrifos air sample data example

```
/*– – – – – – – – – – – – – – – – – – – – – – – – – – – –
SUBJECTID – – applicator (random effect)
    X – – treated at least one enclosed crawl space
    or not (0 = no, 1 = yes)
```

```
    Y – – chlorpyrifos air concentration
    LOD – – limit of detection calculated on the
    chlorpyrifos air concentration
------------------------*/
/*-------LOD/2 METHOD --------*/
DATA air2;
  SET air;
  IF C=1 THEN LY=LOG(LOD/2);
  IF C=0 THEN LY=LOG(Y);
  RUN;
PROC MIXED DATA=air2 COVTEST;
  CLASS SUBJECTID;
  MODEL LY = X/solution cl;
  RANDOM SUBJECTID;
  RUN;
/*--------MLE METHOD --------*/
DATA air3;
  SET air;
  IF C=1 THEN NY=LOD;
  IF C=0 THEN NY=Y;
  RUN;
PROC NLMIXED DATA=air3;
  PARMS SB_2=1 SW_2=1 BETA0=1 BETA1=0;
  BOUNDS SB_2 SW_2>0;
  PI=2*ARSIN(1);
  MU=BETA0 + BETA1*X + U_I;
  IF C=0 THEN L=(1/(SQRT(2*PI*SW_2)*NY))
  *EXP(-(LOG(NY)-MU)**2/(2*SW_2));
  IF C=1 THEN L=PROBNORM((LOG(NY)-MU)/
  SQRT(SW_2));
  LL=LOG(L);
  MODEL NY ~ GENERAL(LL);
  RANDOM U_I ~ NORMAL(0,SB_2) SUBJECT=
  SUBJECTID;
  RUN;
```

Program B:

SAS code for the two-group censored repeated measures model with different within-worker variances

Suppose SAS dataset ONE contains a variable identifying the subjects (SUBJECTID), an indicator variable for group (G: 0, 1), the observed exposure measurement (Y), an indicator variable for censoring (C: 0, 1), and the limit of detection (LOD). The SAS code to generate maximum likelihood estimates of the variance components (common between subject variance σ_B^2; group 0 within-subject variance σ_0^2, and group 1 within-subject variance σ_1^2) and the fixed effects (β_0 and β_1) is given by

```
DATA TWO;
  SET ONE;
  IF C=1 THEN NY=LOD;
  IF C=0 THEN NY=Y;
  RUN;
PROC NLMIXED DATA=TWO;
  PARMS BETA0=0 BETA1=0 SB_2=2 SW0_2=
  2 SW1_2=2;
  BOUNDS SB_2 SW0_2 SW1_2>0;
  MU=BETA0 + BETA1*G + U_I;
  PI=2*ARSIN(1);
  IF G=0 AND C=0 THEN L=(1/(SQRT(2*PI*SW0
  _2)*NY))*EXP(-(LOG(NY)-MU)**2/(2*SW0_2));
  IF G=0 AND C=1 THEN L=PROBNORM
  ((LOG(NY)-MU)/SQRT(SW0_2));
  IF G=1 AND C=0 THEN L=(1/(SQRT(2*PI*
  SW1_2)*NY))*EXP(-(LOG(NY)-MU)**2/
  (2*SW1_2));
  IF G=1 AND C=1 THEN L=PROBNORM
  ((LOG(NY)-MU)/SQRT(SW1_2));
  LL=LOG(L);
  MODEL NY ~ GENERAL(LL);
  RANDOM U_I~NORMAL(0,SB_2) SUBJECT=
  SUBJECTID;
  RUN;
```

Estimates of σ_B^2, σ_0^2, σ_1^2, β_0, and β_1 are is given by SB_2, SW0_2, SW1_2, BETA0, and BETA1, respectively.

Program C:

SAS code for the random intercept and slope censored repeated measures model

Suppose SAS dataset ONE contained a continuous covariate X, and variables for SUBJECTID, Y, C, and LOD.

```
DATA TWO;
  SET ONE;
  IF C=1 THEN NY=LOD;
  IF C=0 THEN NY=Y;
  RUN;
PROC NLMIXED DATA=TWO;
  BOUNDS S1_2 S2_2 SW_2 > 0;
  MU=BETA0 + BETA1*X + A_I + B_I*X;
  IF C=0 THEN L=(1/(SQRT(2*PI*SW_2)*NY))
  *EXP(-(LOG(NY)-MU)**2/(2*SW_2));
  IF C=1 THEN L=PROBNORM((LOG(NY)-MU)/
  SQRT(SW_2));
  LL=LOG(L);
  MODEL NY~GENERAL(LL);
  RANDOM A_I B_I ~ NORMAL([0, 0], [S1_2,
  S12, S2_2]) SUBJECT=SUBJECTID;
  RUN;
```

Here BETA0 represents the intercept, BETA1 represents the slope, S1_2 represents the variance of the random intercepts, S2_2 represents the variance of the random slopes, S12 represents the covariance

between the intercept and slope, and SW_2 represents the within-subject variance.

Program D:

SAS code for the nested random effects model using simulated data

Here, we generate nested random effects data and use NLMIXED to analyze the uncensored data and the 30% censored data.

```
/* GENERATE DATA */
%LET NA = 20;
%LET NB = 3;
%LET NR = 3;
DATA NESTED;
 DO A = 1 TO &NA;
  ERR1 = 3*RANNOR(0);
  DO B = 1 TO &NB;
   ERR2 = 2*RANNOR(0);
   DO REP = 1 TO &NR;
    ERR3 = 1*RANNOR(0);
    RESPONSE = 10+ERR1+ERR2+ERR3;
    OUTPUT;
    END;
   END;
 END;
RUN;
/* ANALYZE UNCENSORED DATA */
PROC NLMIXED DATA=NESTED METHOD=
FIRO;
 ARRAY AEFFECT { 1 };
 ARRAY BEFFECT {&NB };
 MEAN=INTERCEPT + AEFFECT{1} +
 BEFFECT{B};
 RANDOM AEFFECT1 BEFFECT1 BEFFECT2
 BEFFECT3 ~
            NORMAL([0,0,0,0],[VARA,
                              0, VARB_A,
                              0, 0, VARB_A,
                              0, 0, 0, VARB_A])
                              SUBJECT=A;
 *comment: this random statement only works for
NB=3;
 MODEL RESPONSE ~ NORMAL (MEAN,S2);
 RUN;
/* DETERMINE THRESHOLD FOR 30%
CENSORING */
%LET CENSORING=30;
PROC UNIVARIATE DATA=NESTED NOPRINT;
 VAR RESPONSE;
 OUTPUT OUT=P_&CENSORING. PCTLPTS=
 &CENSORING. PCTLPRE=P_;
 RUN;
DATA _NULL_; SET P_&CENSORING.;
 CALL SYMPUT ('LOD', P_&CENSORING.);
 RUN;
/* ANALYZE CENSORED DATA */
DATA NESTED2;
 SET NESTED;
 RESPONSE2=RESPONSE;
 LOD=&LOD.;
 IF RESPONSE<LOD THEN DO;
  C=1;
  RESPONSE2=LOD;
  END;
 ELSE C=0;
 RUN;
PROC NLMIXED DATA=NESTED2;
 ARRAY AEFFECT { 1 };
 ARRAY BEFFECT {&NB };
 PARMS INTERCEPT=10 VARA=9 VARB_A=
 4 S2=1;
 BOUNDS VARA VARB_A S2>0;
 MEAN=INTERCEPT + AEFFECT{1} +
 BEFFECT{B};
 RANDOM AEFFECT1 BEFFECT1 BEFFECT2
 BEFFECT3 ~
            NORMAL([0,0,0,0],[VARA,
                  0 , VARB_A,
                  0 , 0 , VARB_A,
                  0 , 0 , 0 , VARB_A])
                   SUBJECT=A;
 *comment: this random statement only works for
 NB=3;
 PI=2*ARSIN(1);
 IF C=0 THEN L=(1/SQRT(2*PI*S2))*EXP(-
 (RESPONSE2-MEAN)**2/(2*S2));
 IF C=1 THEN L=PROBNORM((RESPONSE2-
 MEAN)/SQRT(S2));
 LL=LOG(L);
 MODEL RESPONSE2 ~ GENERAL(LL);
 RUN;
```

REFERENCES

Amemiya T. (1973) Regression analysis when the dependent variable is truncated normal. Econometrica; 41: 997–1016.

Amemiya T. (1984) Tobit models – a survey. J Econometrics; 24: 3–61.

Burstyn I, Teschke K. (1999) Studying the determinants of exposure: a review of methods. Am Ind Hyg Assoc J; 60: 57–72.

Checkoway H, Pearce N, Kriebel D. (2004) Research methods in occupational epidemiology. 2nd edn. New York, NY: Oxford University Press. pp. 214–6.

Finkelstein MM, Verma DK. (2001) Exposure estimation in the presence of nondetectable values: another look. Am Ind Hyg Assoc J; 62: 195–8.

Flynn MR. (2010) Analysis of censored exposure data by constrained maximization of the Shapiro-Wilk W statistic. Ann Occup Hyg; 54: 263–71.

Gillespie BW, Chen Q, Reichert H *et al*. (2010) Estimating population distributions when some data are below a limit of detection by using a reverse Kaplan-Meier estimator. Epidemiology; 21: S64–70.

Helsel DR. (1990) Less than obvious: statistical treatment of data below the detection limit. Environ Sci Technol; 24: 1766–74.

Helsel DR. (2006) Fabricating data: how substituting values for nondetects can ruin results, and what can be done about it. Chemosphere; 65: 2434–9.

Helsel DR. (2010) Much ado about next to nothing: incorporating nondetects in science. Ann Occup Hyg; 54: 257–62.

Hewett P, Ganser GH. (2007) A comparison of several methods for analyzing censored data. Ann Occup Hyg; 51: 611–32.

Hines CJ, Deddens JA. (2001) Determinants of chlorpyrifos exposures and urinary 3,5,6-trichloro-2-pyridinol levels among termiticide applicators. Ann Occup Hyg; 45: 309–21.

Hornung RW, Reed LD. (1990) Estimation of average concentration in the presence of nondetectable values. Appl Occup Environ Hyg; 5: 46–51.

Hughes JP. (1999) Mixed effects models with censored data with application to HIV RNA levels. Biometrics; 55: 625–9.

Krishnamoorthy K, Mallick A, Mathew T. (2009) Model-based imputation approach for data analysis in the presence of non-detects. Ann Occup Hyg; 53: 249–63.

Kromhout H, Symanski E, Rappaport SM. (1993) A comprehensive evaluation of within- and between-worker components of occupational exposure to chemical agents. Ann Occup Hyg; 37: 253–70.

Laird NM, Ware JH. (1982) Random-effects models for longitudinal data. Biometrics; 38: 963–74.

Lavoué J, Beaudry C, Goyer N *et al*. (2005) Investigation of determinants of past and current exposures to formaldehyde in the reconstituted wood panel industry in Quebec. Ann Occup Hyg 49: 587–602.

Lawless JF. (2003) Statistical models and methods for lifetime data. 2nd edn. New York, NY: John Wiley and Sons. pp. 34–37.

Leidel NA, Busch KA, Lynch JR. (1977) Occupational exposure sampling strategy manual (DHEW [NIOSH] publication no. 77-173). Cincinnati, OH: National Institute for Occupational Safety and Health.

Littell RC, Milliken GA, Stroup WW *et al*. (2006) SAS for mixed models. 2nd edn. Cary, NC: SAS Institute Inc.; pp. 587–9.

McCullagh P, Nelder JA. (1989) Generalized linear models. 2nd edn. New York, NY: Chapman and Hall.

Milliken GA, Johnson DE. (1992) Analysis of messy data, volume 1: designed experiments. New York, NY: Chapman and Hall. pp. 255.

Ogden TL. (2010) Handling results below the level of detection. Ann Occup Hyg; 54: 255–6.

Pettitt AN. (1986) Censored observations, repeated measures and mixed effect models: an approach using the EM algorithm and normal errors. Biometrika; 73: 635–43.

Rappaport SM, Lyles RH, Kupper LL. (1995) An exposure-assessment strategy accounting for within- and between-worker sources of variability. Ann Occup Hyg; 39: 469–95.

SAS Institute Inc. (2004) SAS/STAT® 9.1 User's Guide. Cary, NC: SAS Institute Inc.

Slymen DJ, de Peyster A, Donohoe RR. (1994) Hypothesis testing with values below detection limit in environmental studies. Environ Sci Technol; 28: 898–902.

Swallow WH, Monahan JF. (1984) Monte Carlo comparison of ANOVA, MIVQUE, REML, and ML estimators of variance components. Technometrics; 26: 47–57.

Thiébaut R, Guedj J, Jacqmin-Gadda H *et al*. (2006) Estimation of dynamical model parameters taking into account undetectable marker values. BMC Med Res Methodol; 6: 38.

Thiébaut R, Jacqmin-Gadda H. (2004) Mixed models for longitudinal left-censored repeated measures. Comput Methods Programs Biomed; 74: 255–60.

van der Vaart AW. (2000) Asymptotic statistics (Cambridge series in statistical and probabilistic mathematics). Cambridge, UK: Cambridge University Press; pp. 61–67.