Journal of Biopharmaceutical Statistics
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/lbps20

ROBUST ESTIMATION OF THE MEDIAN LETHAL DOSE

G. E. Kelly [a] & J. K. Lindsey [b]
[a] Department of Statistics, University College Dublin, Belfield, Dublin, 4, Ireland
[b] Department of Biostatistics, Limburgs Universitair Centrum, Diepenbeek, Belgium
Version of record first published: 05 Oct 2011.

To cite this article: G. E. Kelly & J. K. Lindsey (2002): ROBUST ESTIMATION OF THE MEDIAN LETHAL DOSE, Journal of Biopharmaceutical Statistics, 12:2, 137-147

To link to this article: http://dx.doi.org/10.1081/BIP-120014416

MARCEL DEKKER, INC. • 270 MADISON AVENUE • NEW YORK, NY 10016

JOURNAL OF BIOPHARMACEUTICAL STATISTICS
Vol. 12, No. 2, pp. 137–147, 2002

ROBUST ESTIMATION OF THE MEDIAN LETHAL DOSE

G. E. Kelly[1,*] and J. K. Lindsey[2,*]

[1]Department of Statistics, University College Dublin, Belfield, Dublin 4, Ireland
[2]Department of Biostatistics, Limburgs Universitair Centrum, Diepenbeek, Belgium

ABSTRACT

Alternatives to M-estimation for robust estimation of the median lethal dose in biological assays are developed. A class of link functions based on the Student-*t* distribution is proposed, where degrees of freedom are estimated from the data by maximum likelihood. Other alternatives include slash and finite mixture distributions. For bioassays from a pharmaceutical company, these methods extend the standard probit and logistic models, as well as the Huber's M-estimator. They are also applied to several standard examples from the literature.

Key Words: Bioassay; LD_{50}; Robust link function

INTRODUCTION

Biological assays are frequently used by pharmaceutical companies in order to assess the toxicity of certain compounds; see for example Ref. [1]. The present study is motivated by routine assays of a biological compound, to determine its

*E-mail: gabrielle.kelly@ucd.ie and james.lindsey@luc.ac.be

DOI: 10.1081/BIP-120014416
1054-3406 (Print); 1520-5711 (Online)

toxicity, carried out by a pharmaceutical company in Ireland, for which we present data in "Examples" later.

In such a biological assay, doses of a compound at different levels are administered to groups of animals and their responses are recorded. These are usually dichotomized as death (0) or survival (1). The LD_{50} is defined as the unknown dose level for which the probability of response equals one-half. The tolerance of an individual animal is defined to be the dose just insufficient to cause its death. The toxicity of the compound is often summarized in terms of the LD_{50}. However, the assay often has to be repeated because either the estimated LD_{50} or the associated confidence limits do not satisfy regulatory requirements. Because this is expensive, it is appropriate to look at the estimation techniques to find ways in which they could be improved.

The classical approach to the estimation of the LD_{50} is to relate the probability of response to dose using a normal or logistic distribution function as the link function, giving the probit and logistic regression models. The assumed distribution function then represents the population of tolerances and the LD_{50} is estimated by maximum likelihood (MLE). We also note that the problem of estimating an ED_{50}, which occurs in drug screening, is the same as that of the LD_{50}.

In many cases, the form of the tolerance distribution may not be known. This has led some authors to consider a nonparametric approach such as the Spearman–Kärber estimator which is a discretized estimate of the mean of the tolerance distribution. However, in Ref. [2], Miller and Halpern show that the Spearman–Kärber estimator performs poorly for heavy-tailed distributions. These authors proposed a robust approach, giving extensions of L- and M-estimators in the location problem for the LD_{50} in discretized form. In particular, they compared the trimmed mean, the Tukey biweight, and the Spearman–Kärber estimator with regard to the size of their asymptotic variances. Their robust estimators first require an estimate of the sample tolerance distribution, which involves monotonizing the proportion-of-deaths sequence at each dose. This can be problematic, as can the choice of trimming or tuning constants. Moreover, no method for choosing among the estimators has been proposed.

The application that motivated our consideration of heavy-tailed distributions involves routine biological assays on mice. From repeated assays, it has become evident that a tolerance distribution with heavier tails than normal or logistic provides an improved fit to the data. The need for heavy-tails in the tolerance distribution might be interpreted as indicating that some animals are either very frail or very resistant to the compound. In a bioassay, this then prevents the proportion of deaths from being identically one for large doses or zero for small doses.

We, thus, extend traditional approaches and propose using a class of several robust distributions to model tolerances. The best model from the extended class can be chosen using Akaike's information criterion. Pawitan[3] shows how Huber's M-estimator of location can equivalently be defined using a density with a normal center and exponential tails. We consider this interpretation of the Huber estimate as well as the distributions with even heavier tails than Huber.

These include the Student-*t*, slash, or finite mixture distributions, the latter to be defined later. The Student-*t* is appealing because it can model both heavy-tailed and normal tolerance distributions.

ROBUST ESTIMATION

Let $x_1, \ldots, x_k$ be a sequence of dose levels at which subjects are tested. At each dose level, n_i subjects are tested. Define the random variable Y by $Y = 1$ if the outcome for a subject is death and $Y = 0$ if the outcome is survival, and let $\pi = P(Y = 1)$. The response Y will vary with dose. The probit model assumes that

$$\pi = \Phi(\alpha + \beta x) \tag{1}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function and x is either the dose or the log dose. Equation (1) implies that the median lethal dose, LD_{50}, is given by $\rho = -\alpha/\beta$.

The distribution function, here $\Phi(\cdot)$, is known as the tolerance distribution. As already mentioned, the tolerance of a mouse is the dose just insufficient to cause its death. This biological interpretation of the model leads to its appeal in bioassays. Reference [4] contains many examples of its use. Model (1) assumes that the tolerances have a symmetric distribution that is normal. Of course, distributions other than the normal can be considered.

Thus, in general, we shall denote the tolerance distribution by $F(\cdot)$. We assume that all subjects, whether tested at the same dose or different doses, act independently. Thus, the number of responses y_i at dose x_i is binomial with parameters n_i and $\pi_i = F(\alpha + \beta x_i)$. For m different doses, the likelihood of the data is given by

$$L(\alpha, \beta) = \prod_{i=1}^{m} \binom{n}{y_i} [F(\alpha + \beta x_i)]^{y_i} [1 - F(\alpha + \beta x_i)]^{n_i - y_i}.$$

Numerical optimization methods can be used to find the MLE's $\hat{\alpha}$ and $\hat{\beta}$. Their variances and covariances can be estimated from the inverse of the sample information matrix. The MLE of the LD_{50} is $\hat{\rho} = -\hat{\alpha}/\hat{\beta}$. The model can be reparameterized in terms of ρ in order to obtain its standard error (SE) directly. In the case where x is log dose, direct calculation of the SE requires parameterization in terms of $\exp(\rho)$. The asymptotic normal approximation will be more appropriate on one scale than on the other, as will be indicated by Akaike's information criterion.

Alternative Tolerance Distributions

The slash distribution, used in Ref. [2], is the distribution of a normal variable divided by a uniform [0,1] variable. It has heavier tails than the Cauchy

and has been used in Monte Carlo studies of robust estimators for the case of i.i.d. observations. It has the distribution function given by

$$F(x) = \int_0^1 \Phi(ux)\,\mathrm{d}u$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal variable.

For the Huber estimator, consider the one-sample location setting. Huber's $\psi_k(\cdot)$ is given by

$$\psi_k(y, \theta) = \begin{cases} y - \theta & |y - \theta| \le k \\ +k & y > \theta + k \\ -k & y < \theta - k \end{cases}$$

and the corresponding estimator of θ is the solution of $\sum_{i=1}^{m} \psi_k(y_i, \theta) = 0$. By analogy with MLE, the function ψ_k replaces the derivative of the log likelihood, yielding an M-estimator.

As in Ref. [3], $\sum_{i=1}^{m} \psi_k(y_i, \theta) = 0$ can be viewed as a score equation. The associated log of the hypothetical tolerance distribution L_i is, for fixed k, that of a normal center with exponential tails given by

$$\frac{\delta \log(\mathrm{L_i})}{\partial \theta} = \psi(y_i, \theta) \Rightarrow \log(L_i) = \int \psi_k(y_i, \theta)\mathrm{d}\,\theta$$

$$= \begin{cases} -(y_i - \theta)^2/2 & |y_i - \theta| \le k \\ -k(2|y_i - \theta| - k)/2 & |y_i - \theta| > k \end{cases}$$

We normalize this class, so they correspond to distribution functions and choose to estimate k by maximizing this class. When the tolerance distribution is chosen from this class (suitably renormalized), we call it Huber. We note that because of the normalizing constant these distributions no longer have derivative of the log likelihood equal to $\psi_k(\cdot)$.

We also consider a mixture distribution for handling cases in which animals are very frail or very resistant. This assigns probability ν_1 to group responses of 0%, probability ν_2 to group responses of 100%, and probability $(1 - \nu_1 - \nu_2)$ to the responses that are generated from a logistic tolerance distribution (but this could also be assumed to be normal or some other distribution). An example of a dose–response function for this distribution can be seen later.

Tolerance distributions such as the Cauchy and the slash will perform better than the normal if the underlying tolerance distribution is very heavy-tailed. The Student-t distribution provides a wider, alternative model, with members lying somewhere between the normal and the Cauchy.

We use Akaike's information criterion to select the best model. In our examples, this amounts to choosing the model which minimizes twice the negative

log likelihood, with an extra penalty of two units per parameter for the link functions containing more than two unknown parameters, Huber, Student-*t*, and finite mixture.

To obtain likelihood ratio intervals for the LD_{50}, we use profile likelihood: we proceed as in Ref. [1] but perform a sequential grid search for the endpoints. At each step, we use the estimates from the previous step as our new initial estimates. The range of the grid search is taken from the Wald interval based on the SE of the LD_{50}.

EXAMPLES

Pharmaceutical Assay

A pharmaceutical company in Ireland carries out a biological assay on mice on a regular basis. Doses of a compound (a toxin produced by bacteria) at eight levels are administered to the mice, with 10 mice at each dose level. The death or survival of each mouse over 24 hr is the outcome measure. The eight log-doses used by the company bracket the target $\log(LD_{50})$ of 2 and are equally spaced between 1.57 and 2.44. These log-doses are in fact dilution factors, so that a small log-dose will result in a high response; see Ref. [5]. In over 50% of recent assays, only seven doses were administered; some of the responses are shown in Table 1. The standard method is then to estimate the LD_{50} using a probit model with log (dose) rather than dose in Eq. (1). In performing these analyses, it has become clear that the normal model is not always the most appropriate one.

The results for different models fitted to the three assays are given in Table 2, where the 95% confidence interval is the likelihood ratio interval. In Ref. [6], the authors used Newton's method to compute likelihood ratio intervals for the LD_{50}. As described earlier, we performed a sequential grid search to overcome this problem.

In terms of Akaike's information criterion, the mixture is penalized for using four fitted parameters, the Student-*t* and Huber for three parameters, and the logistic, normal, Cauchy, and slash for two. As expected, the Cauchy, slash, and mixture distributions fitted better than the normal or logistic in all three cases.

Table 1. Responses from Three Biological Assays

	Log Dose							
Study	1.57	1.69	1.82	1.94	2.07	2.19	2.31	2.44
(a)	9	9	10	4	1	0	0	—
(b)	—	9	10	10	5	0	0	0
(c)	10	9	10	9	7	0	0	—

Table 2. Results from Three Biological Assays, in Order from Table 1

Distribution	$\widehat{\mathrm{LD}}_{50}$ (SE)	$-\log(\hat{L})$	95% C.I.
Study (a)			
Logistic	82.6 (6.00)	20.84	(71.5, 95.6)
Normal	80.9 (6.13)	21.42	(69.8, 94.5)
Huber	81.2 (6.03)	21.37	(69.9, 94.2)
Cauchy	85.6 (3.34)	18.98	(76.0, 93.7)
Slash	85.3 (3.89)	19.01	(76.1, 94.4)
Mixture	89.2 (5.41)	17.23	(78.9, 102.6)
Student-*t*	87.4 (—)	18.69	(76.8, 93.1)
Study (b)			
Logistic	113.0 (7.05)	15.77	(99.8, 128.3)
Normal	110.3 (7.74)	17.15	(96.3, 127.4)
Huber	120.0 (6.76)	15.17	(104.8, 129.1)
Cauchy	116.3 (1.40)	12.63	(111.9, 121.0)
Slash	116.2 (1.77)	12.63	(111.5, 121.3)
Mixture	116.9 (6.97)	10.18	(108.5, 132.5)
Student-*t*	116.3 (—)	12.30	(109.4, 122.7)
Study (c)			
Logistic	116.5 (7.83)	17.78	(101.9, 133.7)
Normal	113.5 (8.37)	18.86	(98.6, 132.7)
Huber	120.0 (6.75)	17.48	(98.7, 132.4)
Cauchy	119.7 (4.01)	15.10	(113.2, 134.7)
Slash	120.3 (4.47)	15.13	(112.4, 134.6)
Mixture	121.2 (—)	12.62	(114.4, 143.4)

The mixture may be preferred slightly because it allows for greater asymmetry in the response distribution. The Huber distribution has log likelihood between that of the normal and the Cauchy, but it contains an extra parameter.

For all three assays, the mixture distribution with parameters ν_1 and ν_2 converged to a mixture with parameter ν_2 only and thus tended to give the largest estimates of the LD_{50}. The estimate of the parameter k in the Huber model was about 2.0 for the three assays. In the first two assays, the estimated numbers of degrees of freedom for the Student-*t* were 0.38 and 0.49, respectively, giving results close to the Cauchy. In the third assay, the estimated number of degrees of freedom was so small that it converged to the Cauchy.

A plot of the observed frequencies and expected values under the logistic, mixture, and Student-*t* models for assay (a) are shown in Fig. 1. It can be seen that the mixture and Student-*t* fits follow the data more closely than the logistic.

Some regulations require that the estimated LD_{50} be in the interval (80, 120) and that the confidence limits be in the range 68–138; see "Discussion." For the first assay the lower limit of 68 is almost reached by the normal and Huber. For the third assay the slash and mixture have estimated LD_{50} above 120. Thus, the choice of tolerance distribution may be critical in determining whether regulations are satisfied.

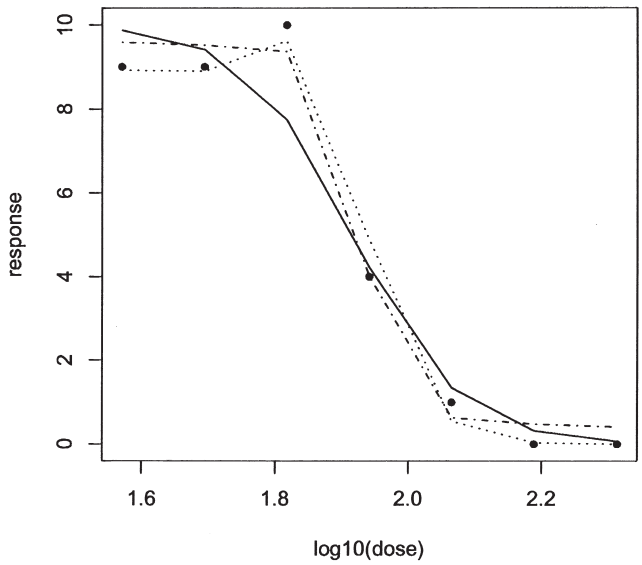

Figure 1. A plot of observed and expected values under the logistic model (—), the mixture model ($\cdots$), and the Student-t ($-\cdot-$) for the data of study (a) of Table 1.

Inhalation Bioassay

The LD_{50} experiments within a Bayesian framework are considered in Ref. [7]. Table 1 of their paper reports inhalation test data with five animals at each dose, where the dose levels in mg/ml are (422, 744, 948, 2069) and the corresponding numbers of deaths are $(0, 1, 3, 5)$. These data have also been considered in Refs. [1,8]. Racine et al.[7] found that the probit model applied to these data gave fiducial limits for the LD_{50} covering the whole real line. This led them to a Bayesian analysis with an uninformative prior for α and β. However, they use log (dose) rather than dose in the probit model. Here we analyze the data using log (dose) as the covariate, with the results given in the top panel of Table 3.

The fits of the logistic and the mixture were the same because the mixture gave estimates of both mixing parameters close to zero. Our fit for the logistic model agrees with Refs. [1,8]. We see, in terms of $\log(\hat{L})$, that a normal distribution fit better than the Cauchy and thus it is a matter of choosing the correct model for the likelihood approach to give reasonable answers. Because the estimated number of degrees of freedom for the Student-t model was 303, its fit is similar to that of the normal model. The Huber model was also similar to the normal with the parameter k estimate of 6.1. The results of all models were quite similar indicating no extremes in the data. They are close to the Bayesian analysis of Ref. [7]. Robust models like Huber's or the Student-t give the same answer as traditional ones when the data are well behaved.

On the original scale, the Student-t distribution gave the results shown in the bottom panel of Table 3, which are similar to those in the top panel. However,

Table 3. Results for the Inhalation Biological Assay

Distribution	$\widehat{\text{LD}}_{50}$ (SE)	$-\log(\hat{L})$	95% C.I.
Log scale			
Logistic	895.3 (83.59)	5.89	(728, 1224)
Normal	895.9 (89.28)	5.87	(723, 1236)
Huber	895.9 (89.28)	5.87	(723, 1235)
Cauchy	912.1 (71.21)	6.33	(739, 1195)
Slash	903.6 (77.10)	6.29	(736, 1217)
Mixture	895.3 (83.59)	5.89	(728, 1224)
Student-*t*	895.8 (89.11)	5.87	(723, 1236)
Original scale			
Logistic	900.7 (77.01)	5.92	(746, 1270)
Normal	899.6 (80.19)	5.89	(746, 1288)
Cauchy	916.9 (64.06)	6.32	(746, 1212)
Slash	909.5 (70.68)	6.29	(746, 1241)
Student-*t*	899.6 (80.18)	5.89	(747, 1287)
Mixture	900.7 (77.01)	5.92	(746, 1270)

Source: Racine et al. (1986).[7]

the estimated number of degrees of freedom was only 27.7. A normal distribution again provided the best fit. For the Huber model, the MLE of the parameter k was ∞, indicating a normal distribution. The mixture tolerance distribution again gave the same results as the logistic.

In another example, Ref. [7] presented hypothetical toxicity data. The doses in mg/ml were (50, 200, 300, 400, 2000), the numbers of animals at each dose were (5, 10, 5, 10, 5), and the numbers of deaths were (1, 4, 2, 6, 4). We analyzed these data using both log (dose) and dose as the covariate, with the results given in Table 4.

The estimated number of degrees of freedom for the Student-*t* distribution was 0.61 using log (dose), providing the best fit to the data although all models with log (dose) were close. Using dose, the estimated Student-*t* degrees of freedom was only 0.26. Using Akaike's criterion, the logistic, normal, Cauchy, and slash fitted equally well and gave similar estimates. Here we see, in the lower panel of Table 4, that the models with dose fitted much more poorly than with log-dose apart from the mixture. The estimates of the mixing parameters were both 0.01 for log-dose whereas for dose they were approximately equal also at 0.11. Thus the mixture on the original scale gave the smallest estimate of the LD_{50} of all models but did not give the best fit by Akaike's criterion. This example illustrates that a number of alternative models should always be tried.

The probit model, using log (dose), gave a much shorter likelihood ratio interval than that reported by Ref. [8] using fiducial limits. However, using dose, the likelihood is flat for some distributions near the maximum so that the MLE of the LD_{50} is given as an interval. The lower confidence limit does not exist for

Table 4. Results for the Hypothetical Biological Assay

Distribution	$\widehat{\mathrm{LD}}_{50}$ (SE)	$-\log(\hat{L})$	95% C.I.
Log scale			
Logistic	317.1 (145.68)	21.98	(48, 5116)
Normal	317.8 (147.56)	21.98	(47, 5309)
Huber	287.2 (180.99)	22.15	(18, 37951)
Cauchy	312.6 (126.45)	21.95	(54, 4179)
Slash	314.3 (136.65)	21.96	(52, 4577)
Mixture	294.4 (556.50)	21.97	(—, 5111)
Student-*t*	311.7 (113.19)	21.93	(72, 4186)
Original scale			
Logistic	520.9–522.9 (324.96)	22.58	entire real line
Normal	535.9–538.3 (—)	22.59	entire real line
Huber	293.5–294.4 (270.71)	22.35	(—, 1380)
Cauchy	385.2 (245.93)	22.42	entire real line
Slash	430.7 (350.66)	22.50	entire real line
Mixture	227.7 (122.32)	21.93	(—, 1118)
Student-*t*	337.2 (94.84)	22.02	(—, 5295)

Source: Racine et al. (1986).[7]

the Student-*t* and Huber distributions; the Cauchy, slash, normal, and logistic models have intervals covering the entire real line.

Reference [1] proves that useful limits based on Fieller's theorem only exist if the Wald test on the slope is significant. At some α level, insignificance will occur with the Wald test and the corresponding confidence interval will be infinite. We note that the 95% Fieller confidence limits did not exist for any model apart from the Student-*t*. The Bayes interval of Ref. [7] was (74.3, 1524). This is considerably shorter than any in Table 4 apart from Huber's. This is because it conditions on a positive slope, i.e., $P(\beta) > 0 = 1$ a priori.

DISCUSSION

Regulatory requirements regarding pharmaceutical assays are complex. For example, in the USA and the European Union, regulations require that the estimated LD_{50} of a vial of the compound with nominal LD_{50} of 100, be in the interval (80, 120). This is subsequent to the estimated LD_{50} having first been normalized with respect to two standards. In addition, a monotone sequence of at least four responses must be obtained and the estimate calculated by probit analysis from this monotone sequence. Regulations in the UK not only carry a similar requirement for a monotone dose sequence but also have requirements concerning the associated Fieller confidence limits. These limits must be in the range 68–138. The present work is part of a larger study to derive an acceptable

alternative to these regulatory requirements which would result in fewer animals being challenged.

As can be seen from Table 2, one important result of fitting a more appropriate tolerance distribution to data from biological assays is that the confidence intervals for the LD_{50} estimate may be much more narrow than those obtained with the logistic and probit link functions. This is important in reducing costs, in satisfying regulatory requirements, and in preventing extensive experimentation on animals. This, in turn, promotes animal rights.

Fitting the Student-*t* distribution provides a useful benchmark in bioassay analysis, in that the estimated number of degrees of freedom, when very small, indicates that a Cauchy or a slash link function is appropriate. When very large, this indicates that a probit or logistic link function should be used. The models considered here need not be tried and compared in practice.

Apart from our mixture model, we have assumed that the tolerance distribution is symmetric, which may not always be reasonable. The proportions of frail and resistant animals for the tested compound may not be the same. The standard link function based on an asymmetric distribution function is the complementary log function. In certain situations, it may be useful to develop robust, heavy-tailed extensions, as we have done for the logistic and probit link functions, perhaps based on asymmetric members of the stable family.

It is not common to consider robust estimation in the way we have, i.e., by obtaining MLEs based on heavy-tailed distributions. Our approach considers a larger class of models than traditional approaches. We then choose among the models using Akaike's information criterion. Our approach has the advantage that our estimation procedures are more readily programmed in standard software packages than classical robust estimators, such as M-estimators with Huber's $\psi(\cdot)$ function or the Tukey biweight.

The approach in our examples, choosing the model which maximizes $\log(\hat{L})$, and using Akaike's information criterion, is widely accepted within the likelihood framework. This is unlike the Bayesian approach,[7] where the question of which prior to use has no universal answer. See Ref. [9] for a more extensive discussion of these different approaches to inference.

ACKNOWLEDGMENTS

We wish to thank R. Gentleman and R. Ihaka,[10] and the R core group for developing the R software that is freely available under the GNU license. Our models were fitted with the probability distribution functions plogis, pnorm, pcauchy, and pt for the Student-*t*. The MLE estimations for the models were performed using the function nlm. The slash distribution function was found by integration using the int function from the second author's R public library called rmutil available at http://www.luc.ac.be/~jlindsey/rcode.html. The models can also be fitted directly using the function gnlr from the library gnlm available as above.

REFERENCES

1. Grieve, A.P. On Likelihood and Bayesian Methods for Interval Estimation of the LD_{50}. In *Statistics in Toxicology*; Morgan, B.J.T., Ed.; Oxford University Press: Oxford, 1996; 87–100.
2. Miller, R.; Halpern, J. Robust Estimators for Quantal Bioassay. Biometrika **1980**, *67*, 103–110.
3. Pawitan, Y. *All Likelihood*; Oxford University Press: Oxford, 2001.
4. Finney, D.J. Estimation of the Median Effective Dose. *Probit Analysis*, 3rd Ed.; Cambridge University Press: London, 1971; 20–42.
5. Kelly, G.E. The Median Lethal Dose-Design and Estimation. J. R. Stat. Soc., Ser. D **2001**, *50*, 41–50.
6. Alho, J.M.; Valtonen, E. Interval Estimation of Inverse Dose-Response. Biometrics **1995**, *51*, 491–501.
7. Racine, A.; Grieve, A.P.; Fluhler, H.; Smith, A.F.M. Bayesian Methods in Practice: Experiences in the Pharmaceutical Industry (with Discussion). Appl. Stat. **1986**, *35*, 93–120.
8. Aitkin, M. Statistical Modelling: The Likelihood Approach. J. R. Stat. Soc., Ser. D **1986**, *35*, 103–113.
9. Lindsey, J.K. Some Statistical Heresies. J. R. Stat. Soc., Ser. D **1999**, *48*, 1–40.
10. Ihaka, R.; Gentleman, R. R: A Language for Data Analysis and Graphics. J. Comput. Graphical Stat. **1996**, *5*, 299–314.

Received December 2001
Revised April 2002
Accepted May 2002