PHARMACEUTICAL STATISTICS
Pharmaceut. Statist. 2009; **8**: 279–286
Published online 25 September 2008 in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/pst.348

Bayesian fitting of a logistic dose–response curve with numerically derived priors

MAIN PAPER

L. W. Huson*,† and N. Kinnersley
[1]*Biostatistics Group, F. Hoffman La-Roche, Welwyn, UK*

In this report we describe the Bayesian analysis of a logistic dose–response curve in a Phase I study, and we present two simple and intuitive numerical approaches to construction of prior probability distributions for the model parameters. We combine these priors with the expert prior opinion and compare the results of the analyses with those obtained from the use of alternative prior formulations.

Keywords: ***Bayesian; simulation; bootstrap***

1. INTRODUCTION

There is no consensus in the literature on what constitutes the best approach to formation of priors when conducting a Bayesian analysis of a logistic dose–response model – a number of different techniques have been described, each involving a different construction for the prior distributions for the model parameters. Many authors assign independent uniform or normal priors for the parameters of the model [e.g. 1–5]. In recent unpublished work, Gelman *et al.* [6] suggest the use of Cauchy priors. Gossl and Kuhkenhoff [7] use flat priors, and Al-Awadhi and Garthwaite [8] construct multivariate normal priors. Hardwick *et al.* [9] use a prior proportional to a likelihood derived using prior information and also incorporate a factor to allow for the degree-of-belief in the strength of the prior (see also [10]). Chen *et al.* [11] assume that data from an earlier study are available in order to provide information for formation of a prior, which does not have a simple closed form, but which can readily be used in Monte Carlo Markov Chain (MCMC) computations.

One difficulty is that direct elicitation of prior opinion about the parameters of a logistic model is unrealistic, as experts are unlikely to be able to say anything directly about the values of model parameters, even if they are knowledgeable in general about the relationship being modelled. It is therefore considered preferable to focus on elicitation of 'observable quantities' and then to construct priors for the parameters themselves from this indirect prior information [12,13]. The earliest approach along these lines was described

*Correspondence to: L. W. Huson, Biostatistics Group, F. Hoffman La Roche, Falcon Way, Welwyn Garden City, UK AL7 1TW .
†E-mail: les.huson@roche.com

Pharmaceutical STATISTICS

by Tstutakawa [14], in which prior information on response probabilities at two specific chosen covariate levels is used to construct a prior for the model parameters themselves (see also [12,15]). This approach was generalized by Bedrick *et al.* [13,16], who use the change-of-variable technique to construct priors for a broader class of linear models, following elicitation of information on response values at a number of levels of a covariate vector.

In this report, we describe two simple numerical approaches for formation of prior probability distributions for the parameters of a logistic dose–response curve, and we apply these techniques to expert prior opinion gathered prior to the initiation of a Phase I dose–response study. We compare our numerically derived priors with a Tsutakawa prior also derived from the expert prior opinion.

2. ELICITATION PROCESS

In our Phase I study, we planned to investigate the rate of occurrence of a specific type of cardiovascular adverse event in each of six cohorts of patients. Each cohort was planned to contain two placebo-treated patients and six patients receiving the test treatment, with ascending doses of a candidate new drug given to each cohort. To this experimental data, we planned to fit a conventional two-parameter logistic dose–response model, relating the probability p of the adverse event to the dose administered: $p = \exp(a+b^*\text{dose})/(1+\exp(a+b^*\text{dose}))$. We needed to form prior distributions for the two parameters a and b, which we refer to below as the 'intercept' and the 'dose-effect' parameters, respectively.

We consulted a group of five internal experts. To begin the elicitation process, we first met with the experts and provided them with a description of the basic ideas of Bayesian statistical methods, and then described our plans for elicitation of prior information for our Phase I study. The primary elicitation technique that we used was to provide each expert with a blank graphical display, showing on the y-axis the probability of toxicity, and on the x-axis dose levels of the test compound. We asked each expert to mark a point on the graph for each dose level, indicating his/her best estimate of the probability of toxicity at that dose. We also asked each expert to provide an estimate of a toxicity probability for each dose, which was plausible, but which they were '90% confident' they could rule out. Experts were not asked to provide these estimates in any particular order or sequence. As a check on consistency of this process, we also asked each expert to give written numerical values – in tabular form – of toxicity probabilities for a specified range of doses, and also to give written numerical estimates for doses at which they would expect to see a pre-specified set of toxicity probabilities. All of these elicited values, in both graphical and tabular form, were checked for internal consistency and plausibility. In this report we make use of the probability estimates elicited graphically, though these were very similar to the tabular estimates.

Our representation of prior belief therefore consisted of two sets of probability estimates provided by each of five internal experts. Each set consisted of an estimated toxicity probability on placebo and at each of the six dose levels that were to be tested in our Phase I study. The first set represented the expert's 'best' estimate of the toxicity probability (which we treat as an estimate of the median of a distribution describing the expert's prior belief) and the second set contained the expert's estimates of the 90th percentile of the toxicity probability at each dose.

3. DERIVING PRIOR DISTRIBUTIONS

We wished specifically to develop a simple programmatic method of forming priors for the parameters of our logistic dose–response curve. This permits the technique to be generalized for use with more complex dose–response models should the logistic model prove to be inadequate. We adopted two approaches to this, both of which

DOI: 10.1002/pst

are simple and intuitive, and neither of which appears to have previously been used in this context.

3.1. CDF technique with simulation

The first method was based on fitting of beta distributions, combined with simulation, and proceeds as follows:

Step 1: For each expert and each dose level, we interpreted the two numerical probability values provided ('best' and '90%') as the median and 90th percentile of a beta distribution, which distribution we therefore assume adequately describes the expert's prior belief. Using these two quantiles, it is a simple process to fit a beta distribution to describe the expert's prior belief about the toxicity probability at a given dose, and this technique is widely used in the context of describing prior beliefs about individual probabilities (see e.g. [17,18]). This procedure was applied to each expert's estimates at each dose, to give a set of beta distributions representing the prior beliefs about the toxicity probabilities. Figure 1 shows an example of the information elicited from one of our experts. In this figure, the expert's median and 90th percentiles are shown for each dose, together with the beta distributions fitted to each of these sets of quantiles.

Step 2: Having fitted a beta distribution at each dose, for each expert, we then simulated random values from these fitted distributions, and to each

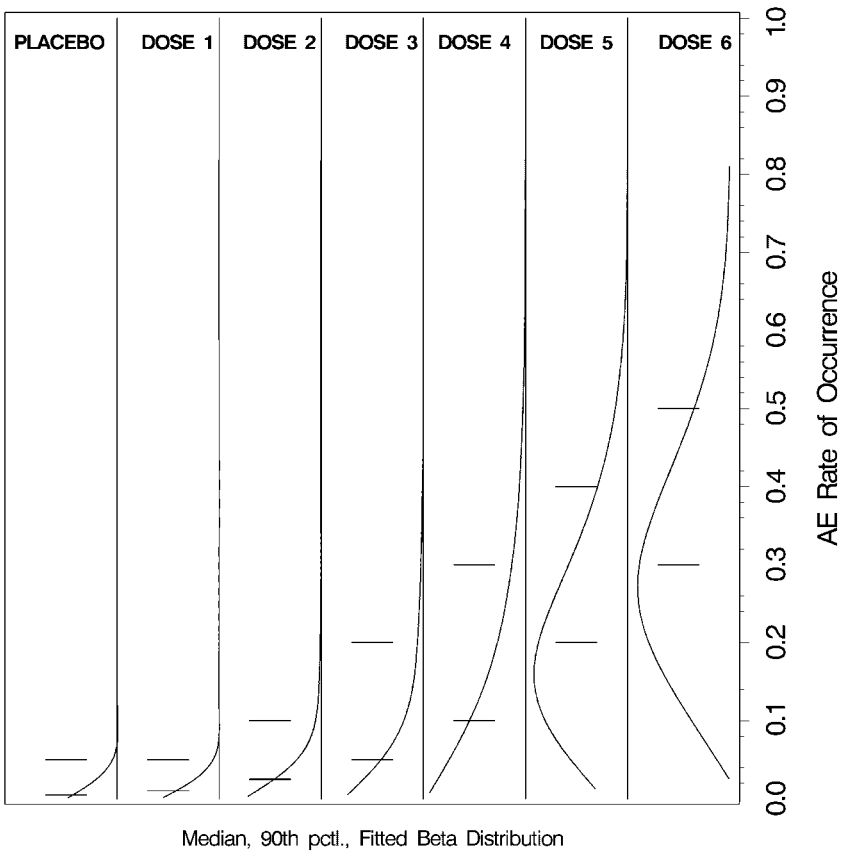

Figure 1. An example of one expert's elicited medians and 90th percentiles, together with fitted beta distributions.

DOI: 10.1002/pst

Pharmaceutical STATISTICS

set of simulated data we fitted a two-parameter logistic dose–response curve. For each such fitting, we stored the values of the two parameters of the fitted curve, and repeated the process to build-up empirical distributions for the two parameters. The repeated fitting was performed separately for each expert's prior opinion, and then the empirical results pooled across all experts. We formed these empirical distributions by generating 2000 simulated data sets.

Step 3: To the empirical parameter distributions built up by this process of simulation we fitted re-scaled beta distributions, which we then use as the priors for the two parameters in the Bayesian fitting of the final logistic dose–response curve.

3.2. Bootstrapping

The second technique we used is based on drawing simple bootstrap samples directly from the expert prior opinion:

Step 1: We drew repeated bootstrap samples of probabilities and doses directly from the expert prior information, pooling across all experts (hence treating the estimates as though they were real data). For the bootstrap sample, we drew 90% of the values from the experts 'best' estimates of the toxicity probabilities, and 10% from the 90th percentiles, to give appropriate weightings to each set of estimates.

Step 2: To each bootstrap sample we fitted a two-parameter logistic curve, and by repeating this process 2000 times we constructed empirical distributions for each of the two parameters.

Step 3: We fitted re-scaled beta distributions to these empirical parameter distributions, and used the fitted beta distributions as priors in the Bayesian fitting of a logistic dose–response curve.

3.3. Checking for stability

With most problems involving simulation and bootstrapping, the appropriate number of samples to draw is best assessed empirically in the context of a particular application. We adopted this approach by increasing the sample sizes in our simulation and bootstrapping processes until the fitted re-scaled beta distributions stabilized. We found little change in the parameters of the fitted re-scaled beta distributions beyond 1000 samples in simulation, and beyond 500 samples for the bootstrap approach.

The beta distributions derived from these two numerical approaches are shown in Figure 2, and for comparison we also show Tsutakawa priors derived using the same expert opinion.

Figure 2 shows marginal prior distributions, but in fact the two parameters in the fitted logistic curves are correlated. This could be taken into account in the Bayesian fitting process, although in

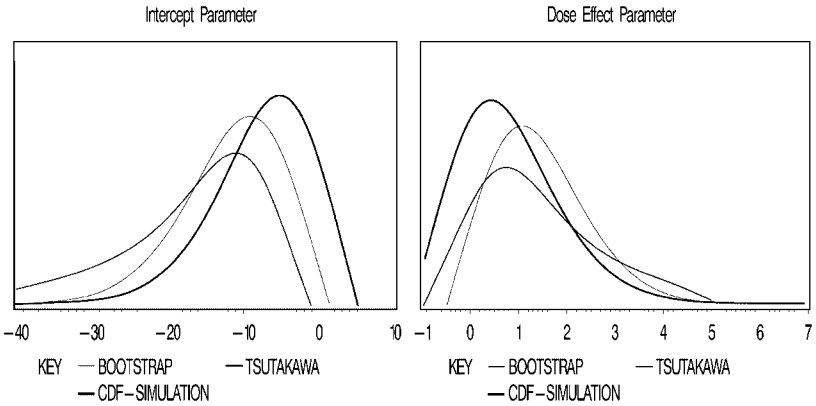

Figure 2. Tsutakawa prior and re-scaled beta distributions fitted to empirically derived priors.

DOI: 10.1002/pst

many previously published analyses authors have assumed independence of the priors, and in our application the results we obtained were similar whether we modelled this correlation or assumed independence.

4. RESULTS OF FITTING OF THE LOGISTIC DOSE–RESPONSE CURVE TO PHASE I DOSE–RESPONSE DATA

As we were performing a purely exploratory analysis of our Phase I study, we used WINBUGS [19] to approximate the posterior distributions for our parameters using the priors described above.

Table I. Posterior parameter estimates for logistic dose–response models using different priors.

Mean and SD of posterior distributions	Intercept	Dose effect
Bootstrap-beta priors	−15.43 (4.746)	2.12 (0.6846)
Tsutakawa prior	−13.71 (7.496)	1.00 (1.1260)
CDF/simulation beta priors	−15.29 (5.385)	2.11 (0.7955)
Bootstrap–independent beta priors	−14.21 (4.368)	1.95 (0.6502)
CDF/simulation–independent beta priors	−12.63 (4.376)	1.72 (0.6556)

WINBUGS models with our beta priors are easily fitted and converge rapidly and successfully. Estimates from the posterior distributions are summarized in Table I.

We also show in Figure 3 a graphical summary of the prior and posterior distributions derived from the analysis using the CDF-simulation technique. Also shown in this plot are the maximum likelihood parameter estimates derived from the experimental data alone. For both the parameters, it can be seen that the posterior distribution, as expected, has moved towards the estimates suggested by the data, and that the posterior means are closer to the values estimated from data alone than to the priors.

For the purposes of our exploratory analysis, the various sets of posterior estimates – both means and standard deviations – are broadly similar, though the Tsutakawa prior gives standard deviations slightly larger than those from our beta priors. The plotted logistic curves are also very similar. There is little difference between the posterior estimates derived from bootstrapped and CDF-simulation priors in our particular application. In further exploratory analyses, we also applied other forms of prior to our data (including normal and uniform priors with various parameters), and found that the results we obtained were similar for most of the priors we tried.

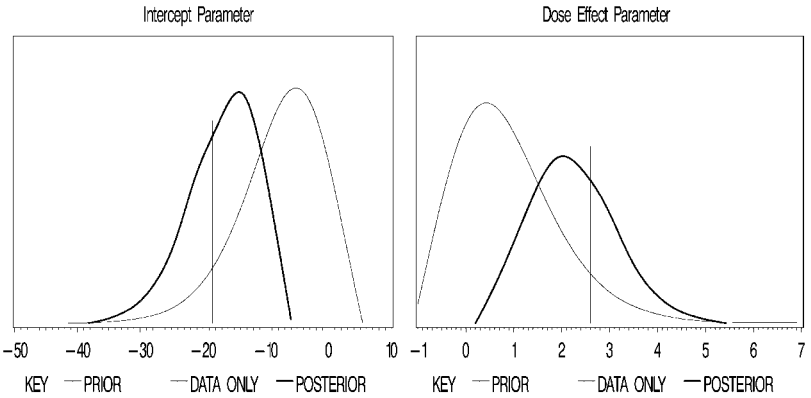

Figure 3. Prior and posterior distributions for the intercept and dose-effect parameters, as estimated using the CDF-simulation method, plus estimates made from experimental data alone.

DOI: 10.1002/pst

Pharmaceutical STATISTICS

5. DISCUSSION

The essential feature of our approach is to use numerical techniques applied to elicited expert opinion to build-up empirical distributions for the parameters of the dose–response model. We do this by repeatedly fitting models to samples drawn from the expert prior opinion, the samples being drawn either by simulation or by bootstrapping. Then, having drawn such samples and constructed empirical distributions for the parameter values, we fit distributions to those empirical results, and use these as priors. We selected the re-scaled beta distribution as our final prior form because of its flexibility and simplicity, but use of the beta distribution at this stage is not an essential feature of our approach, and there are clearly many other possibilities at this stage of the process – for example, normal or gamma distributions, or bivariate distributions.

In the case of applications involving the logistic model, a number of alternative formulations of priors have been described in the literature, and our approach adds a further option. However, the most immediate advantage of our technique is that it could equally be applied to arbitrarily more complex models, as well as the logistic. Since the prior opinion elicited in our approach is not model specific, if any lack-of-fit of the chosen model is detected at the final analysis, the analysis can be repeated using the same prior opinion but with a different model form for the dose–response relationship.

Methods for determining priors for a response probability by eliciting quantiles, as we employed in our study, are discussed in detail by O'Hagan *et al*. [17 (p. 124)], citing in particular the work of Winkler [20] on this and other methods. One of the observations made by O'Hagan *et al*. is that prior distributions elicited in this was may be 'too tight' i.e. may not fully reflect the true level of uncertainty in the prior opinion. Some of the experimental evidence for this is discussed by O'Hagan *et al*. One thing that is not clear is the extent to which this possible 'tightness' may be alleviated by the process of fitting of a beta (or any other) distribution to the elicited quantiles. The approach that we employed in this study is reported by O'Hagan *et al*. to be 'not as tight' as some of the published alternatives, and they conclude that 'much more work is needed before firm recommendations can be made as to which is the best method'. Nevertheless, it is important for practitioners to be aware of possible tightness of priors elicited in this way and to consider possible alternatives.

A reviewer of the first version of this report commented on the fact that, in generating simulated values from the beta distributions, which we had fitted, we had not incorporated any correlation into the values simulated at successive doses. We made a specific decision to generate simulated values without correlation, in the expectation that this would generate data with greater spread than would simulation with correlation, and that this might compensate to some extent for the possible tightness of the individual beta priors fitted at each dose. However, the incorporation of such a correlation into the simulation phase of our method is an alternative, which practitioners may wish to consider.

In our study we combined the prior opinions of our experts by a simple linear pooling, but this is only one of a number of possible methods that have been reported in the literature (see [17] (p. 181) for a recent discussion. In selecting this approach, we were guided by the view that such a simple linear pooling is appropriate in situations in which it is desired to give equal weight to all values considered reasonable by *any* expert. Alternative approaches are possible in which more weight is given to the views of some experts, – i.e. the views of some experts are allowed to 'over-rule' the views of others – and this should be considered in situations in which the opinions of some experts are considered more valuable or more reliable than others. After reviewing a number of methods for both mathematical and behavioural aggregation of expert opinion, O'Hagan *et al*. [17] conclude that 'simple aggregation methods work well in comparison with more complex methods'.

The idea of simulation of data as part of a technique for assembling information on prior distributions has clear connections to the long-established process of using Monte Carlo techniques to study the propagation of uncertainty in

DOI: 10.1002/pst

deterministic mathematical models (see e.g. [21]). Direct bootstrapping of expert prior opinion to form priors appears not to have been reported before in the context of a Bayesian analysis, although Meyer and Booker [22] report bootstrapping of expert prior 'data' to summarize information, which is then used directly in policy making. Elsner and Jagger [23] also use bootstrapping to construct a prior from real data.

In conclusion, we found that the process of graphical elicitation of prior belief was well received by the experts we consulted, and that the two options for formation of priors for the actual parameters of the logistic curve were robust and simple to implement. Application of Bayesian methodology in the context of a Phase I study, especially where conclusions will largely be used internally, may be a valuable formal way of using internal expert opinion.

ACKNOWLEDGEMENTS

We thank Professor John Whitehead for his valuable comments on an earlier version of this work following a presentation made at the PSI Annual Conference in May 2007. This was also presented at the Joint Statistical Meetings in Salt Lake City in July 2007. We are also grateful to an anonymous reviewer for valuable comments on the first version of this manuscript.

REFERENCES

1. Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian data analysis* (2nd edn). Chapman and Hall: New York, 2004.
2. Jaakkola T, Jordan M. A variational approach to Bayesian logistic regression models and their extensions. In Proceedings of the Sixth International Workshop on Artificial Intelligence and Statistics, 1997.
3. Mila AL, Yang XB, Carriquiry AL. Bayesian logistic regression of soybean sclerotinia stem rot prevalence in the US north-central region: accounting for uncertainty in parameter estimation. *Phytopathology* 2003; **93**(6):758–764.
4. Liu Y, Johnson WO, Gold EB, Lasley BL. Bayesian analysis of risk factors for anovulation. *Statistics in Medicine* 2004; **23**:1901–1919.
5. Spears FM, Brown BW, Atkinson EN. The effect of incomplete knowledge of parameter values on single and multiple-stage designs for logistic regression. *Biometrics* 1997; **53**(1):1–10.
6. Gelman A, Jakulin A, Grazia Pittaus M. A default prior distribution for logistic and other regression models. Unpublished manuscript: see www.stat.columbia.edu/gelman, 2006.
7. Gossl C, Kuchenhoff H. Bayesian analysis of logistic regression with an unknown change point and covariate measurement error. *Statistics in Medicine* 2001; **20**:3109–3121.
8. Al-Awadhi SA, Garthwaite PH. Quantifying expert opinion for modelling fauna habitat distributions. *Computational Statistics* 2006; **21**(1):121–140.
9. Hardwick J, Meyer MC, Stout QF. Directed walk designs for dose–response problems with competing failure modes. *Biometrics* 2003; **59**:229–236.
10. Meyer MC, Laud PW. Predictive variable selection in generalized linear models. *Journal of the American Statistical Association* 2002; **97**:859–871.
11. Chen M-H, Ibrahim JG, Yiannoutsos C. Prior elicitation, variable selection, and Bayesian computation for logistic regression models. *Journal of The Royal Statistical Society Series B* 1999; **61**(1):223–242.
12. Tsutakawa RK, Lin HY. Bayesian estimation of item response curves. *Psychometrika* 1986; **51**: 251–267.
13. Bedrick EJ, Christensen R, Johnson W. A new perspective on priors for generalized linear models. *Journal of the American Statistical Association* 1996; **91**(436):1450–1460.
14. Tsutakawa RK. Bayesian inference for bioassay. Technical Report No 52, Department of Statistics, University of Missouri, 1975.
15. Grieve AP. A Bayesian approach to the analysis of LD50 experiments. In *Bayesian statistics*, **Vol. 3** Bernardo JM, DeGroot MH, Lindley DV, Smith AFM (eds). Oxford University Press: Oxford, 1988; 617–630.
16. Bedrick EJ, Christensen R, Johnson W. Bayesian binomial regression: predicting survival at a trauma center. *The American Statistician* 1997; **51**:211–218.
17. O'Hagan A, Buck CE, Daneshkhah A, Eiser R, Garthwaite P, Jenkinson D, Oakley J, Rakow T. Uncertain judgements: eliciting experts' probabilities. ISBN-10: 0-470-02999-4, 2006.
18. Van Dorp JR, Mazzuchi TA. Solving for the parameters of a beta distribution under two quantile constraints. *Journal of Statistical Computation and Simulation* 2000; **67**:189–201.
19. Spiegelhalter D, Thomas A, Best N, Lunn D. WinBUGS user manual, version 1.4', MRC Biostatistics Unit, Institute of Public Health and Department of Epidemiology & Public Health, Imperial

DOI: 10.1002/pst

College School of Medicine (http://www.mrc-bsu.cam.ac.uk/bugs), 2005.
20. Winkler RL. The assessment of prior distributions in Bayesian analysis. *Journal of The American Statistical Association* 1967; **62**:776–880.
21. O'Neill RV, Gardner RH, Mankin JB. Analysis of parameter error in a nonlinear model. *Ecological Modelling* 1981; **8**:297–311.
22. Meyer MA, Booker JM. Eliciting and analyzing expert judgment: a practical guide. ASA-SIAM series on statistics and applied probability. Philadelphia, 2001.
23. Elsner JB, Jagger TH. A hierarchical Bayesian approach to seasonal hurricane modeling. *Journal of Climate* 2004; **17**(14):2813–2827.

DOI: 10.1002/pst