Implications of using different types of priors in a three-level hierarchical Bayesian Model for the analysis of Drug Adverse Events.

Stephen J.W. Evans[1], David Prieto-Merino[1], David J. Spiegelhalter[2], John Whittaker[1], Liam Smeeth[1].

[1]Dept. of Epidemiology and Population Health, London School of Hygiene & Tropical Medicine, UK [2]MRC Biostatistics Unit, Institute of Public Health, Cambridge, UK

Issues of harm caused by drugs are a major concern for regulatory authorities and health professionals as well as patients. It is difficult to assess whether an AE (Adverse Event) observed in a particular patient has been caused by the drug or is coincidental. In randomised trials we compare the incidence of each AE in a group of patients taking the drug with a control group not taking the drug, and differences that are statistically significant may be accepted as evidence for causality. A multiplicity problem with a considerable potential for false positive findings appears when many AEs are compared. The same issues arise in analysis of observational studies or in spontaneous reports of adverse events.

Methods of adjustment for multiplicity such as Bonferroni usually treat each AE as independent. This is not sensible if AEs tend to be clinically related, as they typically are: if a drug increases the incidence of a given AE, an increase is expected in the incidence of other related AEs. Mehrotra and Heyse (2004)[1] analysed data from a clinical trial using a frequentist false discovery rate method that took account of grouping AEs into "Body Systems" (BS) and then grouping body systems into one "Body" group (B). Berry and Berry (2004)[2] have suggested an alternative 3 level Bayesian Hierarchical Model (BHM) that they implemented on the Mehrotra and Heyse data.

The main advantage of a BHM is that information can be borrowed between the units when they are in the same group (for instance several AEs in a Body System). The posterior distributions for each AE have two interesting features: 1) their means are shrunk to their group mean; 2) the more homogeneous the observed data of the AEs in a group, the smaller the variance of each posterior. This shrinks related terms towards each other and reduces the variance, making it easier to discard the null hypothesis for some AEs. This agrees with clinical common sense. With a BHM, a crucial decision is the form of the prior distribution. Berry & Berry used a "lump and smear" mixed prior distribution that puts a "spike" of probability on the null, and distributes the rest as a Normal centred on the null.

We have implemented the Berry & Berry model in WinBUGS, and we have found some potentially undesirable features. The spiked priors can produce bimodal posteriors that are difficult to explain clinically. Moreover, these priors represent a strong belief that the precise null is likely, which may not always be appropriate. We believe investigation of sensitivity to these priors is important. For instance, a standard hierarchical model formulation using Gaussian priors avoids bimodality in the posteriors and produces notably different detailed results. The table shows the observed effect of the drug in some AEs and the posterior summaries (Ln(OR) with 95% Intervals) using both priors.

BS,AE	Name	Data	(exact 95%CI)		Post. with Mix-prior			Post. With Gauss. priors		
[3,4]	Diarrhoea	0.86	(0.03	1.75)	0.73	(0.00	1.55)	0.80	(0.13	1.51)
[10,4]	Rash	1.42	(0.09	3.13)	1.29	(0.00	2.31)	1.29	(0.43	2.23)
[10,6]	Rash, rubella-like	1.50	(-0.02	5.81)	1.10	(0.00	2.21)	1.13	(0.20	2.15)

[1] Mehrotra DV, Heyse JF. *Stat Methods Med Res*. 2004 Jun;13(3):227-38.
[2] Berry S & Berry D *Biometrics (*2004) 60:418-26