BUGS

Blocker: random effects meta-analysis of clinical trials

Carlin (1992) considers a Bayesian approach to meta-analysis, and includes the following examples of 22 trials of beta-blockers to prevent mortality after myocardial infarction.

Study	Mortality: deaths / total	
	Treated	*Control*
1	3/38	3/39
2	7/114	14/116
3	5/69	11/93
4	102/1533	127/1520
.....		
20	32/209	40/218
21	27/391	43/364
22	22/680	39/674

In a random effects meta-analysis we assume the true effect (on a log-odds scale) δ_i in a trial *i* is drawn from some population distribution.Let r^C_i denote number of events in the control group in trial *i*, and r^T_i denote events under active treatment in trial *i*. Our model is:

$$r^C_i \sim \text{Binomial}(p^C_i, n^C_i)$$

$$r^T_i \sim \text{Binomial}(p^T_i, n^T_i)$$

$$\text{logit}(p^C_i) = \mu_i$$

$$\text{logit}(p^T_i) = \mu_i + \delta_i$$

$$\delta_i \sim \text{Normal}(d, \tau)$$

"Noninformative" priors are given for the μ_i's and d; two alternative "noninformative" priors are considered for the random effects variance: prior 1 uses a Gamma(0.001, 0.001) prior on the precision τ, while prior 2 assumes a proper uniform prior on the standard deviation σ. The graph for this model (with prior 2) on is shown below. We want to make inferences about the population effect d, and the predictive distribution for the effect δ_{new} in a new trial. *Empirical Bayes* methods estimate d and τ by maximum likelihood and use these estimates to form the predictive distribution $p(\delta_{new} \mid d_{hat}, \tau_{hat})$. *Full Bayes* allows for the uncertainty concerning d and τ.

Graphical model for blocker example (with prior 2):

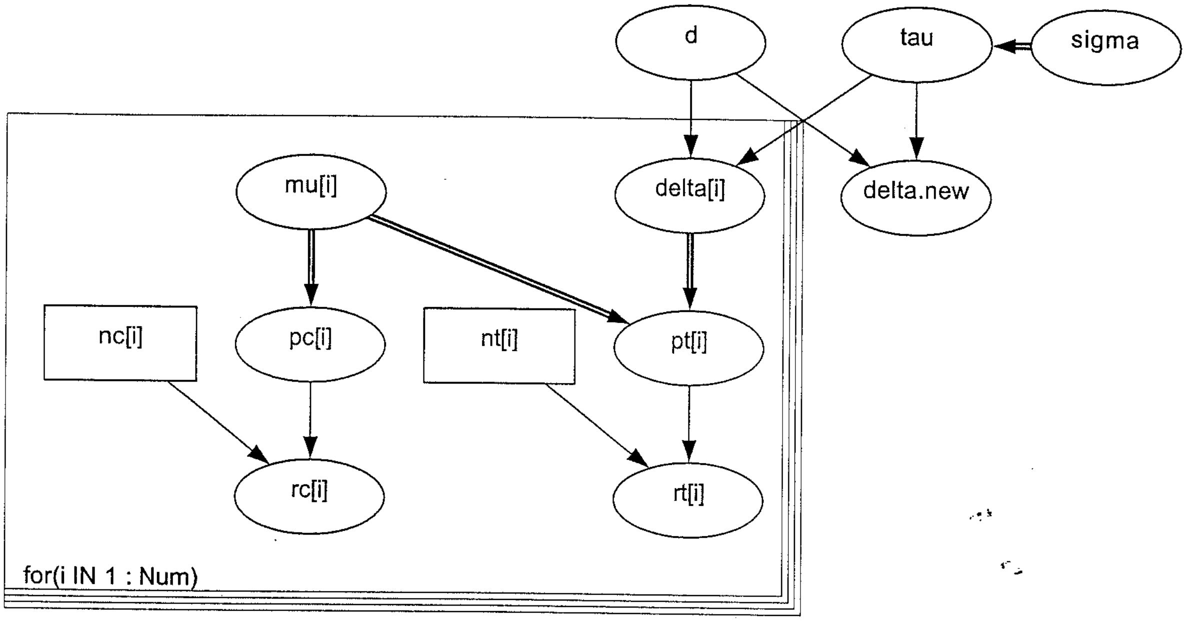

BUGS *language for blocker example:*

```
model
{
  for( i in 1 : Num ) {
     rc[i] ~ dbin(pc[i], nc[i])
   rt[i] ~ dbin(pt[i], nt[i])
   logit(pc[i]) <- mu[i]
   logit(pt[i]) <- mu[i] + delta[i]
   mu[i] ~ dnorm(0.0,1.0E-5)
   delta[i] ~ dnorm(d, tau)
  }
  d ~ dnorm(0.0,1.0E-6)
  # Choice of priors for random effects variance
  #tau ~ dgamma(0.001,0.001)
  #sigma <- 1 / sqrt(tau)
  tau<-1/(sigma*sigma)
  sigma~dunif(0,10)
  delta.new ~ dnorm(d, tau)
}
```

Data ⇨

```
list(rt = c(3,  7,  5,  102,  28, 4,  98,  60, 25, 138, 64, 45,  9, 57, 25, 33, 28, 8, 6, 32, 27, 22 ),
   nt = c(38, 114, 69, 1533, 355, 59, 945, 632, 278,1916, 873, 263, 291, 858, 154, 207, 251, 151, 174, 209, 391, 680),
   rc = c(3, 14, 11, 127, 27, 6, 152, 48, 37, 188, 52, 47, 16, 45, 31, 38, 12, 6, 3, 40, 43, 39),
   nc = c(39, 116, 93, 1520, 365, 52, 939, 471, 282, 1921, 583, 266, 293, 883, 147, 213, 122, 154, 134, 218, 364, 674),
Num = 22)⇦
```

Inits2 ⇨

```
list(d = 0, delta.new = 0, sigma=1, mu = c(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0),
    delta = c(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)) ⇦
```

Inits ⇨

```
list(d = 0, delta.new = 0, tau=1, mu = c(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0),
    delta = c(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)) ⇦
```

Results

A 1000 update burn in followed by a further 10000 updates gave the parameter estimates using the gamma prior on τ

node	mean	sd	MC error	2.5%	median	97.5%	start	sample
d	-0.2489	0.06282	0.002297	-0.3734	-0.248	-0.1239	1001	10000
delta.new	-0.2496	0.1576	0.002582	-0.5773	-0.2514	0.07974	1001	10000
sigma	0.1243	0.06834	0.002835	0.02878	0.1142	0.2796	1001	10000

Our estimates are lower and with tighter precision - in fact similar to the values obtained by Carlin for the empirical Bayes estimator. The discrepancy appears to be due to Carlin's use of a uniform prior for σ^2 in his analysis, which will lead to increased posterior mean and standard deviation for d, as compared to our use of a gamma(0.001, 0.001) prior on the precision which is approximately equivalent to assuming $p(\sigma^2) \sim 1 / \sigma^2$ (see his Figure 1).

If we use a uniform prior on σ, the estimate of σ is slightly increased but there is little influence on the overall conclusions.

node	mean	sd	MC error	2.5%	median	97.5%	start	sample
d	-0.2484	0.06517	0.002568	-0.3699	-0.2493	-0.1161	1001	10000
delta.new	-0.2489	0.1692	0.002666	-0.5963	-0.2532	0.1073	1001	10000
sigma	0.1334	0.07934	0.00502	0.009759	0.1263	0.305	1001	10000

In some circumstances it might be reasonable to assume that the population distribution has heavier tails, for example a t distribution with low degrees of freedom. This is easily accomplished in *BUGS* by using the dt distribution function instead of dnorm for δ and δ_{new}.