


Name: Solutions

1. Draw a directed graph of this model.

5



2. Copy the line or lines of WinBUGS code that specifies/specify the second stage of the model.

$$p[i] \sim \text{dbeta}(a, b)$$

3. Are the outcomes for all the individual baby rats in all litters considered exchangeable?

Briefly explain.

No. The rats within each litter are considered exchangeable, but not rats in different litters. Each rat has to contribute to its litter's x_i so its data is associated with the correct p_i .

~

8

4. Write an expression to which the joint posterior density of all the model parameters is proportional. Show your work.

χ

$$\prod_{i=1}^L \left[\binom{n_i}{x_i} p_i^{x_i} (1-p_i)^{n_i-x_i} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p_i^{a-1} (1-p_i)^{b-1} \right] a^{L-1} e^{-0.2a} b^{L-1} e^{-0.2b}$$

\mathcal{L}

$$\prod_{i=1}^L \left[p_i^{x_i+a-1} (1-p_i)^{n_i-x_i+b-1} \right] \frac{[\Gamma(a+b)]^L}{[\Gamma(a)\Gamma(b)]^L} e^{-0.2a} e^{-0.2b}$$

5. From the attached WinBUGS output, what are the posterior mean and the 95% equal-tail credible set for the parameter of primary interest (the population proportion of baby rats that would survive 21 days after birth if their mother ingested the chemical during pregnancy)? (numeric answer)

χ

0.7438 (0.597, 0.8514)

6. Explain the Bayesian interpretation of the credible set you provided in the previous question.

χ There is 95% probability that the true mean of the distribution from which litter-specific survival probabilities are drawn, lies in this interval.

7. For the parameter a, the MC error is 0.11 and the estimated post sd is 1.47. A rule of thumb in using MCMC to fit Bayesian models is that the MC error should be less than 1/20th of the posterior sd. What would you do to achieve that with this model and dataset?

χ Run more iterations.

2

8. WinBUGS history plots, autocorrelation plots, and BGR diagnostic plots are given for the parameters a, b, and theta and for one of the p's. (The plots for the remaining p's are similar.) In addition, the table corresponding to the BGR diagnostic plot for a is shown. Use the plots and table to answer the following questions.

(a) How many burn-in iterations would you discard? (Circle the best answer.)

i. none

ii. about 250

iii. about 500

iv. about 750

v. 1000 or more

(b) Briefly justify your answer to the previous question, referring to specific WinBUGS output.

Red line stays below 1.2 from iteration 101 on. Blue and green lines stabilize at around 3.44-3.5 at iteration 251.

9. Give one reason why the hierarchical model shown on page 1 is more appropriate in this application than a model in which the existence of individual litters is ignored and the total number of survivors in all litters is considered a single binomial random variable. (There are many good answers for this.)

Baby rats from the same litter are likely to be similar to each other in survival probabilities due to inherited characteristics. This violates the assumption that a binomial variable is the count of success from independent Bernoulli trials.

10. Suppose that we had data on only the last litter, in which none of the baby rats survived. We wish to use it to infer about the population proportion p . The likelihood is proportional to

$$p^0(1-p)^7$$

We wish to use an improper prior on p , namely

$$p \sim \text{Beta}(0, 0)$$

(a) Is it possible to carry out a valid Bayesian analysis using this prior and the data from litter 16 alone? (Yes/no)

(b) If your answer to the previous question was "Yes," what is the resulting posterior distribution, $p(p|x)$? If your answer to the previous question was "No," why not?

Posterior would be improper!

$$p(p|x_{10}) = \text{Beta}(0+0, 0+7)$$
$$= \text{Beta}(0, 7)$$

↑
not > 0.