A bioassay experiment using the t-link

Gelman *et al.* discuss the following data (in Section 3.7 of the book). To test the toxicity of a new drug, a standard approach is to perform bioassay experiments on animals. Various dose levels of a chemical compound are given to batches of animals and a binary outcome (survivial or death) is recorded for each animal. We consider the following data.

Dose (log g / ml)	Number of deaths	Number of animals
-0.86	0	5
-0.30	1	5
-0.05	3	5
0.73	5	5

The book uses the logistic regression model, but instead of the logistic link we now use the t link. Let y_i be the number of deaths out of n_i at a dose level x_i (now expressed in the units log (g / ml)). We model the outcomes outcomes y_i (number of deaths) as the values of binomially distributed random variables Y_i, which are independent given θ, where

$$[Y_i \mid \theta] \overset{\text{ind}}{\sim} \text{Bin}(n_i, p_i), \qquad i = 1, \ldots, n$$

and the probability p_i is given by

$$p_i = F_\nu(\alpha + \beta x_i),$$

where F_ν is the distribution function of Student's t distribution with ν degrees of freedom. We fix the value of at $\nu = 8$. The parameter vector $\theta = (\alpha, \beta)$, and p_i is a function of θ. We take the flat prior, where α and β are independently normally distributed with large variances.

The logistic regression model is almost the same as our model. You can get the MLE (maximum likelihood estimate) for the logistic model using standard software. In R this goes as follows.

```
x <- c(-0.86, -0.3, -0.05, 0.73)
n <- c(5, 5, 5, 5)
y <- c(0, 1, 3, 5)
response <- cbind(y, n - y)
fit <- glm(response ~ x, family = binomial)
# This is the MLE:
coef(fit)
# And this is its approximate covarince matrix
summary(fit)$cov.unscaled
```

You should get approximately the same results for the posterior expectation and posterior covariance in the Bayesian model, but you should take the implied change of scale into account.

Your task is to explore the posterior distribution using software you write yourself. Now the Gibbs sampler which is outlined in section 8.5 of the lecture notes is a natural choice.

Your tasks are as follows.

- Write a formula for the likelihood.

- Make a contour plot of the posterior density.

- Implement the Gibbs sampler. As a starting point, you will be given an implementation of the Gibbs sampler for the probit regression problem of sec. 8.4.

- Draw a simulated sample from the posterior density on top of the contour plot. Estimate posterior summaries such as the posterior expectation and the posterior covariance matrix.

- A particular parameter of interest is the LD-50, the value of the dose where the probability of death is 50 %. It is given by $\phi = -\alpha/\beta$. Plot a histogram of the posterior distribution of ϕ and calculate a central 95 % credible interval for it.