The following problems are about Bayesian one-way ANOVA with random effects. This model has been discussed by Gelfand and Smith (1990) and many other authors. The hypothesis testing aspect needs Metropolis-Hastings-Green.

The data are denoted

$$y_{ij}, \qquad i = 1, \ldots, k, \quad j = 1, \ldots, n_i$$

and the model is described by

$$y_{ij} \sim \text{Normal}\left(\theta_i, \frac{1}{\lambda_e}\right), \qquad i = 1, \ldots, k, \quad j = 1, \ldots, n_i,$$

and

$$\theta_i \sim \text{Normal}\left(\mu, \frac{1}{\lambda_\theta}\right), \qquad i = 1, \ldots, k.$$

and

$$\begin{aligned}
\mu &\sim \text{Normal}\left(\mu_0, \frac{1}{\lambda_0}\right) \\
\lambda_\theta &\sim \text{Gamma}\left(a_1, b_1\right) \\
\lambda_e &\sim \text{Gamma}\left(a_2, b_2\right)
\end{aligned}$$

The six hyperparameters μ_0, λ_0, a_1, b_1, a_2, and b_2 are assumed to be known, chosen so that the prior distributions represent one's a priori opinion about the parameters. To complete the specification we need to say the obvious, that the y_{ij} are conditionally independent given the θ_i and λ_e and the θ_i are conditionally independent given μ and λ_θ.

We also consider the "null hypothesis" model of "no treatment effect" which sets $\lambda_\theta = \infty$ so the θ_i are concentrated at μ. Essentially this eliminates these variables θ_1, ..., θ_k and λ_θ from the model, leaving only μ and λ_e. In order to do a Bayesian hypothesis test (calculate a Bayes factor) for these two models we need a sampler that samples both as in the section of the notes about Bayesian model comparison.

In order to have some real data, let us use an example problem from Stat 3011 analyzed (frequentistly) on the web at

`http://www.stat.umn.edu/geyer/3011/examp/ch10.html#three-mean-basic`

the data being in the file

`http://rweb.stat.umn.edu/WSdata/Ch10data/reading.txt`

The response (`increase`) is the increase (before and after treatment) on "reading age" as measured by the "Gadapol test" (no idea what it is) the treatments being four methods of teaching reading

Map using diagrams to note and relate main points

Scan reading instructions and skimming for an overview before reading in detail

Both both of the above

Neither neither (the control)

For specificity consider the following hyperparameter values

$$
\begin{aligned}
\mu_0 &= 0 \\
\lambda_0 &= 1 \\
a_1 &= 2 \\
b_1 &= 10 \\
a_2 &= 2 \\
b_2 &= 5
\end{aligned}
$$

In my notation the gamma density is

$$f(x \mid a, b) = \frac{b^a}{\Gamma(a)} x^{a-1} e^{-bx}$$

rather than the other parameterization (with $1/b$ rather than b). This makes all the formulas simpler.

Question 1

Derive Gibbs updates for these two models (the big model with all the parameter and the little or "null hypothesis" model with only μ and λ_e).

Question 2

Derive a Metropolis-Hastings-Green update that jumps between these two models.

Question 3

Run a sampler that computes the Bayes factor for model comparison. Don't forget Monte Carlo standard error for your Bayes factor calculation.