1. Radon (Rmd): Read in the data from the file `mn.radon.data` using the `dget` command. These data consist of radon measurements in Minnesota counties.

   (a) Make a dotplot of the data. Construct the ANOVA table and test for overall differences between counties. Using terms from the ANOVA table, obtain estimates of μ, σ^2 and τ^2, and an estimate of the intraclass correlation coefficient. Describe in words your results.

   (b) Rank order the counties in terms of their sample means, and identify the top 5 and bottom 5 counties. How do the sample sizes of the top and bottom counties compare to the average sample size across groups? Plot county sample means versus county sample sizes, and comment.

   (c) Obtain shrinkage estimates $\hat{\mu}_j$ for each county j using the formula

   $$\hat{\mu}_j = \frac{n_j/\hat{\sigma}^2}{n_j/\hat{\sigma}^2 + 1/\hat{\tau}^2}\bar{y}_j + \frac{1/\hat{\tau}^2}{n_j/\hat{\sigma}^2 + 1/\hat{\tau}^2}\bar{y}_{\cdot\cdot}.$$

   Rank order the counties in terms of the $\hat{\mu}_j$'s, and identify the top and bottom 5 counties. Describe any differences between this ranking and the one in (b).

   (d) Using `lme4`, obtain MLEs for μ, σ^2 and τ^2. Compare them to the estimates obtained above.

   (e) Using the command `mll.oneway` from the course slides, plot the *profile log-likelihood* for each of (μ, τ^2, σ^2) as follows: Fix σ^2 and τ^2 at their MLE values. With these fixed, plot the log-likelihood of μ in a range near the MLE value. Confirm that the peak occurs at the MLE. Make analogous plots for σ^2 and τ^2.

2. Mean squared error (written): In this exercise, we will compute how well $\bar{y}_j$ does at estimating μ_j as compared to the shrinkage estimator $w\bar{y}_j + (1-w)\mu$ in the hierarchical normal model. Here, $w = \frac{n_j/\sigma^2}{n_j/\sigma^2+1/\tau^2}$. To make the calculations simpler, we will assume $\mu = 0$ and $\sigma^2 = 1$, and the value of τ^2 is known, so $w = \frac{n_j}{n_j+1/\tau^2}$ and the shrinkage estimator is just $w\bar{y}_j$.

(a) Compute first the mean squared error (MSE) of $\bar{y}_j$, $\text{MSE}(\bar{y}_j) = \text{E}[(\bar{y}_j - \mu_j)^2]$. Do this by first taking the conditional expectation given μ_j, then the unconditional expectation of the result. The answer should be simple and intuitive.

(b) Now we will compute the MSE of $w\bar{y}_j$, $\text{MSE}(w\bar{y}_j) = \text{E}[(w\bar{y}_j - \mu_j)^2]$. Do this by first computing the conditional expectation $\text{E}[(w\bar{y}_j - \mu_j)^2 | \mu_j]$. One way to do this is by writing $(w\bar{y}_j - \mu_j)^2$ as

$$\begin{aligned}(w\bar{y}_j - \mu_j)^2 &= (w(\bar{y}_j - \mu_j) - (1-w)\mu_j)^2 \\ &= w^2(\bar{y}_j - \mu_j)^2 - 2w(1-w)(\bar{y}_j - \mu_j)\mu_j + (1-w)^2\mu_j^2.\end{aligned}$$

Take the expectation of this quantity, conditional on μ_j (that is, treating μ_j as a constant).

(c) Now take the expectation of the result in (b), assuming $\mu_j \sim N(0, \tau^2)$. This is the MSE of $w\bar{y}_j$. Simplify the result as much as possible.

(d) For what values of n_j and τ^2 is the MSE of $w\bar{y}_j$ lower than that of $\bar{y}_j$? For what values is it much lower?