Homework 3
Due: February 26, 2024

The molecule Nitric Oxide (NO) relaxes blood vessels and improves circulation generally and oxygen flow throughout the brain. It was declared "molecule of the year" by *TIME* magazine in 1992 and the discovery of its biological effects in humans was the basis of the 1998 Nobel Prize in Medicine.

NO is created in the paranasal sinuses and humming transfers NO from the paranasal sinus to the nasal passages, where it can be inhaled and affect circulation. To investigate whether such effects ultimately influence human cognition, a study was run where participants responded as quickly as possible to the appearance a target green circle. Response time was measured, with lower response times indicating better performance. Within a block of 20 trials a participant was asked to

1. Hum through their nose (keeping the mouth closed).
2. Speak, "1, 2, 3, 4…1, 2, 3, 4…"
3. Remain quiet.

If humming had the expected benefits, mean response times would be shortest for the humming condition. The speak condition is an active control, while the quiet condition is a passive control. Theory suggests that the two control conditions should have similar response times. Each participant went through all three conditions (in a random order for each participant).

From the class website, get the CSV file named "HummingRT.csv". For each participant it reports the mean response time (averaged across the 20 trials) for each condition. Use the data to conduct a Bayesian analysis to test for the effects of humming on response time.

Use the *ulam( )* function in the *rethinking* package to build and compare four models:

1. A "null" model that assumes no differences between study conditions but supposes differences between participants.

2. A "full" model that assumes differences between conditions and participants. This should not be a multi-level model.

3. A "predicted benefit" model where the mean for the humming condition must be shorter than the mean for the quiet condition and the mean for the speaking condition must equal the mean for the quiet condition.

4. A "multi-level" model that uses a hyper-parameter to define the prior across participants.

To complete the assignment, submit R source code that generates all these models and provides answers to the following questions. In addition, you should submit a document (MS Word, or PDF) that directly answers the questions.

1. Check the r-hat values for each model parameter. If r-hat is bigger than 1.05, change the *ulam( )* settings and run it again. [Hint: if necessary, increase "iter" and "warmup". You might also want to change your priors to not be so broad.]
2. In a direct comparison of the model WAIC values, which model best predicts future data sets? How much confidence do you have in this evaluation?
3. Extract the coefficients for each participant for the full and multi-level models. Plot them. Does the plot indicate shrinkage for the multi-level model?
4. Does your analysis suggest that humming reduces response times? What happens with the "predicted benefit" model?
5. Write up a "Results" section as if it were part of a journal article.