HW 5

Shravan Vasishth

In these exercises, you can use any or all of the four methods discussed in the lecture. In general, bridgesampling will be time-consuming; Savage-Dickey will be faster.

Exercise 1: Is there evidence for differences in the effect of cloze probability among the subjects?

Use Bayes factor to compare the log cloze probability model that we examined in the section called Non-nested models in chapter 13, with a similar model but that incorporates the strong assumption of no difference between subjects for the effect of cloze ($\tau_{u_2} = 0$).

Exercise 2: Is there evidence for the claim that English subject relative clauses are easier to process than object relative clauses?

Consider again the reading time data coming from Experiment 1 of @grodner presented in exercise @ref(exr:hierarchical-logn):

```
data("df_gg05_rc")
df_gg05_rc
```

```
## # A tibble: 672 x 7
##    subj  item condition    RT residRT qcorrect experiment
##   <int> <int> <chr>     <int>   <dbl>    <int> <chr>
## 1     1     1 objgap      320   -21.4        0 tedrg3
## 2     1     2 subjgap     424    74.7        1 tedrg2
## 3     1     3 objgap      309   -40.3        0 tedrg3
## # i 669 more rows
```

You should use a sum coding for the predictors. Here, object relative clauses (`"objgaps"`) are coded $+1/2$, and subject relative clauses as $-1/2$.

```
df_gg05_rc <- df_gg05_rc %>%
  mutate(c_cond = if_else(condition == "objgap", 1/2, -1/2))
```

Using the bayes_factor function discussed in class, quantify the evidence against the null model (no population-level reading time difference between SRC and ORC) relative to the following alternative models:

a. $\beta \sim Normal(0, 1)$
b. $\beta \sim Normal(0, 0.1)$
c. $\beta \sim Normal(0, 0.01)$
d. $\beta \sim Normal_+(0, 1)$
e. $\beta \sim Normal_+(0, 0.1)$
f. $\beta \sim Normal_+(0, 0.01)$

(A $Normal_+(.)$ prior can be set in `brms` by defining a lower boundary as 0, with the argument `lb = 0`.)

What are the Bayes factors in favor of the alternative models a-f, compared to the null model?

Now carry out a standard frequentist likelihood ratio test using the `anova()` function that is used with the `lmer()` function. The commands for doing this comparison would be:

```
m_full <- lmer(log(RT) ~ c_cond +
                 (c_cond || subj) + (c_cond || item),
               df_gg05_rc)
m_null <- lmer(log(RT) ~ 1 + (c_cond||subj) + (c_cond || item),
               df_gg05_rc)
anova(m_null, m_full)
```

How do the conclusions from the Bayes factor analyses compare with the conclusion we obtain from the frequentist model comparison?

Exercise 3: In the Grodner and Gibson 2005 data, in question-response accuracies, is there evidence for the claim that sentences with subject relative clauses are easier to comprehend?

Consider the question response accuracy of the data of Experiment 1 of Grodner and Gibson 2005.

a. Compare a model that assumes that RC type affects question accuracy on the population level and with the effect varying by-subjects and by-items with *a null model* that assumes that there is no population-level effect present.
b. Compare a model that assumes that RC type affects question accuracy on the population level and with the effect varying by-subjects and by-items with *another null model* that assumes that there is no population-level or group-level effect present, that is no by-subject or by-item effects. What's the meaning of the results of the Bayes factor analysis?

Assume that for the effect of RC on question accuracy, $\beta \sim Normal(0, 0.1)$ is a reasonable prior, and that for all the variance components, the same prior, $\tau \sim Normal_{+}(0, 1)$, is a reasonable prior.