

# Take Home Quiz 3


Your Name

Due Tuesday Nov 26 at 11:00am

There is no time limit, but this quiz should take you about 30-40 minutes. Place your answers into this markdown document, knit it, and hand in the result as a PDF or Word document. You may use R, any reference material, and information already available on the internet. Do not work together, do not use AI, and do not get help from anyone but Dr. Clair.

## Problem 1 (10 points)

Consider the normal regression model  $Y_i \sim N(\mu_i, \sigma)$  with  $\mu_i = \beta_0 + \beta_1 X_i$ . From this plot, give your best estimate (by eye) of  $\beta_0$ ,  $\beta_{0c}$  (centered),  $\beta_1$  and  $\sigma$ .



## Problem 2 (10 points)

In the regression model from Problem 1, suppose we choose priors  $\beta_0 \sim N(200, 100)$  and  $\beta_1 \sim N(0, 40)$ .

- Give an example of any prior plausible line by writing its equation.
- Give an example of any line that would not be prior plausible, by writing its equation.

The last three problems use data on blood alcohol content, available on our web page at <https://turtlegraphics.org/bayes/data/bac.csv>

**Problem 3 (10 points)**

16 students from The Ohio State University were randomly assigned to drink a number of cans of beer. Then their blood alcohol content (BAC) was measured.

Use `rstanarm` to build a normal regression model of `bac` on `beers` using the default weakly informative priors, 4 chains, and 10000 iterations per chain.

Report the mean posterior `beers` coefficient to at least three decimal places. What does this number tell you about beer drinking?

**Problem 4 (10 points)**

What prior did `rstanarm` choose for the  $\sigma$  parameter? What is the mean of the  $\sigma$  prior distribution?

**Problem 5 (10 points)**

Make a density plot of the posterior predictive distribution for a student that drinks 5 beers.

Make a density plot of the posterior predictive distribution for a student that drinks 9 beers.

Both distributions use the same  $\sigma$  values, but one is noticeably wider. Why?