Assignment #8 — Due Wednesday, March 25, 2009, by 5:00 P.M.

Turn in homework in lecture, discussion, or your TA's mailbox. Indicate the discussion section in which you expect to attend to pick up this assignment on the assignment.

311: Monday 1:20–2:10 **312:** Monday 12:05–12:55

This short assignment includes some problems about Bayesian inference and some problems that ask you to review past ideas. Enjoy Spring Break!

1. A sample $s = (24.81, 23.52, 23.58, 23.08, 19.99, 28.98, 24.08, 21.75, 19.23, 23.30)$ is modeled as an i.i.d sample from a Normal(μ, σ^2) distribution.

   (a) Find the sample mean and standard deviation (in R use `mean()` and `sd()`).

   (b) Find a 90% confidence interval for μ using the t-distribution (recall `qt()`).

   (c) Find a 90% confidence interval for μ using the non-parametric bootstrap (recall package `boot` and function `boot()` from a previous assignment).

   (d) Assume that $\sigma = \sigma_0 = 4$, and assume a prior distribution $\mu \sim \text{Normal}(20, 2^2)$. Find the Bayesian posterior distribution for μ and find a 90% credible region for μ.

   (e) Assume a model where both μ and σ^2 are unknown, and let $\nu = 1/\sigma^2$ be the precision of the distribution. Assume that $\nu \sim \Gamma(4, 64)$ and that $\mu \,|\, \nu \sim \text{Normal}(20, 0.2/\nu)$. Find the posterior distribution of (μ, ν).

   (f) You can use simulation to generate an estimated posterior credible region for μ by simulating ν from its posterior distrubution (using `rgamma()`) and then simulating $\mu \,|\, \nu$ from a conditional posterior distribution (using `rnorm()`), and then finding quantiles of the sample (using `quantile()`). Write some R code to do this (following examples from lecture and notes to come).

Work to do, but not turn in.

- Read sections 7.1–7.4, 8.1–2.