STA402L: HOMEWORK 4

DUE: 11:59 PM ON FRIDAY, MARCH 6

Instructions. Solutions must be submitted to Gradescope as a single PDF. Programming exercises must be completed in R, should be clearly presented, and include all R code. Lab questions are restated here for convenience, but you should refer to the lab itself for details.

Total points. Book exercises: 10; Lab exercises 18; Overall: 28.

BOOK EXERCISES

B1. (5 points) Hoff 5.2.

B2. (5 points) Let $Y_1, \ldots, Y_n \overset{\text{iid}}{\sim} \mathcal{N}(\theta, \sigma^2)$ with θ and σ^2 unknown. *Jeffreys' prior* for (θ, σ^2) is

$$\pi(\theta, \sigma^2) \propto (\sigma^2)^{-3/2}.$$

Derive the posterior, $p(\theta, \sigma^2 \mid y_1, \ldots, y_n)$, under this prior. (*Note: You may either integrate directly to verify whether the posterior is proper (not recommended), or rewrite the posterior in a form that allows you to recognize it as a known distribution or combination of distributions. Do all derivations in terms of the variance σ^2 (not the precision). That is, keep the normal likelihood expressed in terms of σ^2 before combining it with the prior.*)

LAB EXERCISES

L1. (2 points) Write down the posterior means of α and β. Give 95% credible intervals for each. Considering the amount of data we have, do the results seem surprising?

L2. (2 points) Compute the posterior means of α and β. Give 95% credible intervals for each. How does the posterior inference under this $N(0, 1)$ prior compare to the diffuse priors above? How informative is this weakly informative prior?

L3. (2 points) Would you say that a Cauchy prior is more or less informative than a Normal prior (assume that their inter-quartile ranges are comparable)?

L4. (2 points) What happens as we increase the number of observations?

L5. (2 points) When might we prefer to use lighter- versus heavier-tailed priors?

L6. (2 points) How might we determine a reasonable scale for the prior?

L7. (2 points) What would we expect to be the posterior mode of our samples? Calculate the posterior mode theoretically, and compare it to the estimated mode from the posterior samples.

L8. (2 points) Is this prior proper, that is, does it integrate to 1? Does it result in a proper posterior, that is, does it integrate to 1? If so, under which conditions?

L9. (2 points) If we set a uniform prior for π, what is the induced prior on $\theta = \frac{\pi}{(1+\pi)}$ (do a quick review of functions of random variables if you do not remember)? Is this prior proper, that is, does it integrate to 1?