Homework 2

*The problems/parts marked * are for practice and won't be graded, so no need to submit solution for those. Of course, recommended working through them*

Problem 1: Suppose $x_1, \dots, x_n \mid \lambda \sim \text{Poisson}(\lambda)$. Consider a gamma prior on λ: $\lambda \sim \text{gamma}(\alpha, \beta)$ with density

$$\frac{\beta^\alpha}{\Gamma(\alpha)} e^{-\beta\lambda} \lambda^{\alpha-1} \mathbb{1}_{(0,\infty)}(\lambda).$$

Verify that

a($*$). The posterior distribution $\lambda \mid x_{1:n} \sim \text{gamma}(T + \alpha, n + \beta)$, where $T = \sum_{i=1}^n x_i$. This in particular implies the posterior mean $E(\lambda \mid x) = (T + \alpha)/(n + \beta) = (T/n) \times n/(n + \beta) + (\alpha/\beta) \times \beta/(n + \beta)$.

b($*$). The marginal likelihood

$$m(x) = [\prod_{i=1}^n x_i!]^{-1} \ \frac{\beta^\alpha}{\Gamma(\alpha)} \times \frac{\Gamma(T+\alpha)}{(n+\beta)^{(T+\alpha)}}.$$

c. Find the predictive distribution of a future observation $x_{n+1} \mid x_{1:n}$ and identify the distribution.

Problem 2: Suppose $x_1, \dots, x_n \mid \mu \sim N(\mu, 1)$. Assume a normal prior $\mu \sim N(\xi, \lambda^{-1})$. Verify that

a($*$). The posterior distribution

$$\mu \mid x_{1:n} \sim N\Big(\frac{n\bar{x} + \xi\lambda}{n+\lambda}, \frac{1}{n+\lambda}\Big).$$

There is a simple way to remember this posterior in normal models with normal priors. The posterior precision (inverse variance) equals the data precision plus the prior precision. The posterior mean is a convex combination of the prior mean and the data mean, with the weights proportional to the respective precisions:

$$E(\mu \mid x_{1:n}) = \frac{n}{n+\lambda}\bar{x} + \frac{\lambda}{n+\lambda}\xi.$$

b. The log-marginal likelihood

$$\log m(x) = -\frac{n}{2}\log(2\pi) - \frac{S}{2} - \frac{1}{2}\ \frac{n\lambda}{n+\lambda}(\bar{x}-\xi)^2 - \frac{1}{2}\log(n+\lambda) + \frac{1}{2}\log(\lambda),$$

where $S = \sum_{i=1}^n)(x_i - \bar{x})^2$. [**Note** that the maximized log-likelihood is $-\frac{n}{2}\log(2\pi) - \frac{S}{2}$.]

c($*$). The predictive distribution:

$$f(x_{n+1} \mid x_{1:n}) \sim N(\hat{\mu}, 1 + \hat{\sigma^2}),$$

where $\hat{\mu}$ and $\hat{\sigma^2}$ are the posterior mean and variance of μ. [Hint: can you use some property of the normal distribution to avoid performing the integral explicitly?]

Problem 3($*$): A random sample of n students is drawn from a large population, and their weights are measured. The average weight of the n sampled students is $\bar{x} = 150$ pounds.

Assume the weight in the population is normally distributed with unknown mean μ and known standard deviation 20 pounds. Suppose your prior distribution for μ is normal with mean 180 and standard deviation 40.

(a) Give a 95% credible interval for μ when $n = 10$. Do the same when $n = 100$. Overlay the two posterior distributions (when $n = 10$ and $n = 100$) with the prior distribution on the same plot.
(b) A new student is sampled at random from the population and has weight x_{new} pounds. Find the posterior predictive distribution for x_{new} when $n = 100$. If we wanted to make a point prediction for x_{new}, what will you suggest? [You can use the results derived in Problem 2.]

Problem 4: Suppose $x \mid p \sim \text{binomial}(n, p)$ and p has a Beta(α, β) prior.

a. Calculate the posterior distribution of $\theta = \text{logit}(p) := \log[p/(1-p)]$.

b. Find a closed form expression for the posterior mean of θ. [Hint: you can get the mean from the moment generating function of the posterior of θ]

c. Suppose $n = 33$ and $x = 20$. Use a Monte Carlo (MC) approach to approximate the posterior mean of θ. Report on the accuracy relative to the number of MC samples.

d. Same data as in part (c). Obtain a 95% credible interval for θ.

Problem 5: Refer to the dataset `normdat.txt` in the homework folder in ecampus. Assume a normal $N(\mu, \tau^{-1})$ model for the data. Consider a NG prior given by $\mu \mid \tau \sim N(0, \lambda^{-1}\tau^{-1})$ and $\tau \sim \text{Gamma}(0.5, 0.5)$. Set $\lambda = 0.1$.

Let $c = \tau^{-1/2}/\mu$ denote the coefficient of variation, which is the ratio of the standard deviation and the mean. Compute a 95% credible interval for c, clearly indicating how you obtain the interval.

Problem 6: A random sample of 1447 adults were interviewed prior to the presidential election in 1988, when 727 favored George Bush, 583 favored Michael Dukakis and 137 had no opinion or favored other candidates. Let θ_1 and θ_2 denote the population proportions in favor of Bush and Dukakis respectively. Assume a uniform prior on the simplex for (θ_1, θ_2) and write down the posterior distribution of (θ_1, θ_2). Calculate (i) the posterior mean and variance of θ_1, (ii) the posterior mean and variance of θ_2, and (iii) the posterior mean of θ_1/θ_2.

Problem 7: Suppose $x_1, \ldots, x_n \mid \mu, \sigma^2 \sim N(\mu, \sigma^2)$. Consider the *improper* prior $\pi(\mu, \sigma^2) \propto 1/\sigma^2$. This is clearly not a proper density, but can be combined with the likelihood in the usual way to yield a proper posterior. Such priors are called *improper priors*. The interest in improper priors stem from the fact that they often provide exactly the same inference as in classical statistics. They can also be interpreted as limits of vague proper priors.

Compute the marginal posterior of μ and obtain a 95% credible interval. Comment on any observations you make.

Problem 8: Consider the regression model $y \mid \beta, \sigma^2 \sim \mathcal{N}_n(X\beta, \sigma^2 \mathrm{I}_n)$ with the multivariate NG prior $\beta \mid \sigma^2 \sim \mathcal{N}_d(0, g\sigma^2(X'X)^{-1})$ and $\sigma^2 \sim \text{IG}(a/2, b/2)$. Find the marginal density of y, integrating over both β and σ^2.

—

If you use more than one sheet of paper, please staple the sheets together. This homework will be due in class on Tuesday Feb 24.