AMS 206B: Intermediate Bayesian Inference
HOMEWORK # 2 (More on Conjugate Families and Decision Theory)

1. An engineer is surveying a large shipment of mechanical parts for quality control and decides to test 10 randomly selected items. Historically, the proportion of defective items has been around 1% and it has rarely been above 2%.

   (a) What would be a reasonable (conjugate) prior distribution for θ (the proportion of defective items), given this historical information? Find the posterior distribution under this prior and the corresponding posterior distribution for a random sample $x_1, \ldots, x_{10}$.

   (b) Assume that the engineer finds no defective components in his survey. What would be the posterior distribution? What would be its mean?

   (c) Find the MLE for θ.

   (d) Which one would you prefer, the MLE or the Bayes estimator? Why?

   (e) If the engineer decides to use a uniform (noninformative?) prior instead, what would be the posterior distribution?

2. Let $X_1, \ldots, X_n$ be an i.i.d. sample such that $X_i|\theta, \sigma^2 \sim N(\theta, \sigma^2)$, where σ^2 is known and μ is unknown. Also, let your prior for θ be a mixture of conjugate priors, i.e.,

$$\pi(\theta) = \sum_{l=1}^{K} w_l \phi(\theta|\mu_l, \tau^2),$$

   where $\phi(\theta|\mu, \tau^2)$ denotes the Gaussian density with mean μ and variance τ^2.

   (a) Find the posterior distribution distribution for θ based on this prior.

   (b) Find the posterior mean.

   (c) Find the prior predictive distribution associated with this model.

   (d) Find the posterior predictive distribution associated with this model.

3. Let $X_1, \ldots, X_n$ be an i.i.d. sample such that $X_i \sim N(\theta, 1)$. Suppose that you know that $\theta > 0$, and you want your prior to reflect that fact. Hence, you decide to set $\pi(\theta)$ to be a normal distribution with mean μ and variance τ^2, truncated to be positive, i.e.,

$$\pi(\theta) = \frac{1}{\sqrt{2\pi}\tau\Phi(\mu/\tau)} \exp\left\{ -\frac{1}{2}\left(\frac{\theta-\mu}{\tau}\right)^2 \right\} \mathbf{I}_{[0,\infty)}(\theta).$$

   (a) Find the posterior distribution distribution for θ based on this prior. Is this a conjugate prior?

   (b) Find the prior predictive distribution.

4. Let $X_1, \ldots, X_n$ be an i.i.d. sample such that each X_i comes from a truncated normal with unknown mean θ and variance 1,

$$f(X_i|\theta) = \frac{1}{\sqrt{2\pi}\Phi(\theta)} \exp\{-\frac{1}{2}(X_i - \theta)^2\}\mathbf{I}_{[0,\infty)}(X_i).$$

If $\theta \sim N(\mu, \tau^2)$, find the posterior for θ. Is this a conjugate prior for this problem? How is this problem different from the previous one?

5. (Robert, reworded to be consistent with our notation). Consider $x \sim Binomial(n, \theta)$ with n known.

   (a) If the prior is $\theta \sim Beta(\sqrt{n}/2, \sqrt{n}/2)$, give the associate posterior.

   (b) What is the estimator that minimizes the Bayesian expected posterior loss if the loss function is $L(\theta, \theta(x)) = (\theta - \theta(x))^2$? Call such estimator $\hat{\theta}(x)$ and show that its associated risk $R(\theta, \hat{\theta}(x))$ is constant.

   (c) Let $\theta_0(x) = x/n$. Find the risk for this estimator, i.e., find $R(\theta, \theta_0(x))$. Compare the risks for $\hat{\theta}(x)$ and $\theta_0(x)$ for $n = 10, 50$, and 100. Conclude about the appeal of $\hat{\theta}(x)$.

6. (Robert, reworded to be consistent with our notation). Consider $x \sim N(\theta, 1)$ and $\theta \sim N(0, n)$. Let $\hat{\theta}(x)$ be the estimator that minimizes the Bayesian expected loss for the square error loss. Show that the Bayesian expected loss evaluated at $\hat{\theta}(x)$ is equal to $n/(n+1)$, i.e., show that $\rho(\pi, \hat{\theta}(x)) = n/(n+1)$.

7. (Adapted from Robert) Consider the LINEX loss function defined by

$$L(\theta, \theta(x)) = e^{c(\theta - \theta(x))} - c(\theta - \theta(x)) - 1.$$

   (a) Show that $L(\theta, \theta(x)) \geq 0$ and plot this loss as a function of $(\theta - \theta(x))$ when $c = 0.1, 0.5, 1, 2$.

   (b) Find $\hat{\theta}(x)$ the estimator that minimizes the Bayesian expected posterior loss.

   (c) Find $\hat{\theta}(x)$ when $x_1, \ldots, x_n \sim N(\theta, 1)$ and $\theta \sim N(\mu, \tau^2)$.

8. Let $L(\theta, \theta(x)) = \omega(\theta)(\theta - \theta(x))^2$, with $\omega(\theta)$ a non-negative function, be the weighted quadratic loss. Show that $\hat{\theta}(x)$, the estimator that minimizes the Bayesian expected loss $\rho(\pi, \theta)$ has the form

$$\hat{\theta}(x) = \frac{E(\omega(\theta)\theta|x)}{E(\omega(\theta)|x)}.$$

Hint: Show that any other estimator has a larger Bayesian expected loss.

9. (Adapted from Robert). Consider $x \sim N(\theta, 1)$, $\theta \sim N(0, 1)$ and the loss

$$L(\theta, \theta(x)) = e^{3\theta^2/4}(\theta - \theta(x))^2.$$

   (a) Show that the estimator that minimizes the Bayesian expected posterior loss in this case is $\hat{\theta}(x) = 2x$. Hint: use the previous exercise.

(b) Show that $\theta_0(x) = x$ dominates $\hat{\theta}(x)$.

10. If $x \sim Binomial(n, \theta)$ and $\theta \sim Beta(\alpha, \beta)$, find the Bayes estimator of θ under loss

$$L(\theta, a) = \frac{(\theta - a)^2}{\theta(1-\theta)}.$$

Note: Be careful about $x = 0$ and $x = n$.

11. Assume you have to guess a secret number θ. You know that θ is an integer. You can perform an experiment that would yield either the number before it or the number after it, with equal probability. You perform the experiment twice. More formally, let x_1 and x_2 be independent observations from

$$f(x = \theta - 1|\theta) = f(x = \theta + 1|\theta) = 1/2.$$

Consider the 0-1 loss function, i.e.,

$$L(\theta, \theta(x)) = \left\{ \begin{array}{ll} 1 & \theta(x) \neq \theta \\ 0 & \text{otherwise} \end{array} \right\}.$$

(a) Find the risks of the estimators $\theta_0(x_1, x_2) = \frac{x_1+x_2}{2}$ and $\theta_1(x_1, x_2) = x_1 + 1$.

(b) Find the estimator $\hat{\theta}(x_1, x_2)$ that minimizes the Bayesian expected loss.

12. Consider a point estimation problem in which you observe $x_1, \ldots, x_n$ as i.i.d. random variables of the Poisson distribution with parameter θ. Assume a squared error loss and a prior of the form $\theta \sim Gamma(\alpha, \beta)$.

(a) Show that the Bayes estimator is $\hat{\theta}(x) = a + b\bar{x}$ where $a > 0$, $b \in (0, 1)$ and $\bar{x} = \sum_{i=1}^{n} x_i/n$. You may use the fact that the distribution of $\sum_i x_i$ is Poisson with parameter θn without proof.

(b) Compute and graph the frequentist risks of $\hat{\theta}(x)$ and that of the MLE.

(c) Compute the Bayes risk of $\hat{\theta}(x)$.

(d) Suppose that an investigator wants to collect a sample that is large enough that the Bayes risk after the experiment is half of the Bayes risk before the experiment. Find that sample size.