1. Suppose that $X$ has a Poisson distribution with unknown mean $\theta$.

   Determine the Jeffreys prior $\pi^J$ for $\theta$, and discuss whether the "scale-invariant" prior $\pi_0(\theta) = 1/\theta$ might be preferrable as noninformative prior.

2. Suppose again that $X$ has a Poisson distribution with unknown mean $\theta$.

   Using $\pi^J$ as the reference measure, find the maximum entropy prior under the constraints that the prior mean and variance of $\theta$ are both 1. (Just write it in terms of the constraints $\lambda_1$ and $\lambda_2$ from lectures, do not solve for these.)

   Repeat, for the reference measure $\pi_0$. (Again, just write it in terms of the constraints $\lambda_1$ and $\lambda_2$ from lectures, do not solve for these.)

3. Suppose that $x_1, \ldots, x_n$ is a random sample from a Poisson distribution with unknown mean $\theta$. Two models for the prior distribution of $\theta$ are contemplated;

$$\pi_1(\theta) = e^{-\theta}, \quad \theta > 0, \text{ and } \pi_2(\theta) = e^{-\theta}\theta, \quad \theta > 0.$$

   (a) Calculate the the Bayes estimator of $\theta$ under both models, with quadratic loss function.

   (b) The prior probabilities of model 1 and model 2 are assessed at probability 1/2 each. Calculate the Bayes factor for $H_0$:*model 1 applies* against $H_1$:*model 2 applies*.

4. * Let $\theta$ be a real-valued parameter and let $f(x|\theta)$ be the probability density function of an observation $x$, given $\theta$. The prior distribution of $\theta$ has a discrete component that gives probability $\beta$ to the point null hypothesis $H_0 : \theta = \theta_0$. The remainder of the distribution is continuous, and conditional on $\theta \neq \theta_0$, its density is $g(\theta)$.

   (a) Derive an expression for $\pi(\theta_0|x)$, the posterior probability of $H_0$.

   (b) Derive the Bayes factor $B(x)$ for the null hypothesis against the alternative.

   (c) Express $\pi(\theta_0|x)$ in terms of $B(x)$.

   (d) Explain how you would use $B(x)$ to construct a most powerful test of size $\alpha$ for $H_0$, against the alternative $H_1 : \theta \neq \theta_0$.

5. Suppose that $x_1, \ldots, x_n$ is a sample from a normal distribution with mean $\theta$ and variance $v$. Let $H_0 : \theta = 0$, and let the alternative be $H_1 : \theta \neq 0$. The prior distribution of $\theta$ has a discrete component that gives probability 1/2 to the point null hypothesis $H_0$; the remainder of the prior distribution is continuous, and conditional on $\theta \neq \theta_0$, its density is $g(\theta)$ given by

$$g(\theta) = (2\pi w^2)^{-1/2} \exp\{-\theta^2/(2w^2)\},$$

   for $-\infty < \theta < \infty$. Show that, if the sample mean is observed to be $10(v/n)^{1/2}$, then

   (a) the likelihood ratio test of size $\alpha = 0.05$ will reject $H_0$ for any value of $n$;

   (b) the posterior probability of $H_0$ converges to 1, as $n \to \infty$.