Estimation and testing exercise*

(1) If objects arrive at random time intervals, with the time between arrivals independent, then if the expected number of arrivals per unit of time (an hour, we'll say) is λ, the pdf of the number of arrivals that actually occur in an hour is

$$p(n \mid \lambda) = \frac{\lambda^n}{n!} e^{-\lambda} \,.$$

This is the **Poisson** distribution. A standard recommendation for a "flat" prior on λ in this problem is to make it flat in $\log \lambda$, i.e. to give it the form $d\lambda/\lambda$.[1] Another possibility is a conjugate prior, which here takes the form of a $\Gamma(q, \alpha)$ distribution, i.e. a pdf

$$\frac{\alpha^q \lambda^{q-1} e^{-\alpha\lambda}}{\Gamma(q)} \,.$$

(a) Show that, if we have an i.i.d. sample n_i, $i = 1, \dots, k$ from this Poisson pdf, then the posterior mean of λ under the $d\lambda/\lambda$ prior is an unbiased estimator in the non-Bayesian sense.

(b) Is there any choice of α and q in the conjugate prior that at the same time makes the prior proper (i.e. makes it integrate to one) and makes the posterior mean classically unbiased?

(c) The unbiased estimator in part (1a) will with positive probability be zero. Explain why no Bayesian estimator derived from a proper prior (other than one that puts probability 1 on $\lambda = 0$) and a symmetric loss function could ever, in any sample, produce an estimate $\hat{\lambda} = 0$.

(d) Find the posterior mean, the posterior median, and a 95% posterior probability interval for λ under the prior pdf $\lambda^{-\frac{1}{2}} e^{-\lambda} / \Gamma(\frac{1}{2})$, assuming a sample of five draws with the n_i given by 5,10,10,5,9. To see how sensitive the results are to the prior, repeat the analysis using the $d\lambda/\lambda$ prior. By using the fact that the posterior has the shape of a Γ distribution, you can do this problem without a computer, if you have access to tables of a Γ or χ^2 distribution.

(e) Are either of the 95% probability intervals you found in (1d) non-Bayesian 95% confidence intervals? Explain your answer. [This is too hard to do as an exercise. It is included here as a point of information. The interval derived from a $d\lambda/\lambda$ prior is in fact a non-Bayesian 95% confidence interval.]

[1] This is the form of the Jeffreys prior in this problem.

(2) You arrive at home at 7PM to find the phone is not working. It was working 10 hours ago when you left. For some decision you need to make, it is important when the phone service shut down. Your answering machine contains a single message: a call from the phone company telling you that later today the phone service will be shut down. The answering machine says this call arrived at 10AM, one hour after you left.

Here is a model for assessing the uncertainty: The time of shutdown is $Y \in (0, 10)$. Conditional on $Y = y$, the time of arrival of the call from the phone company (which we'll call X) is distributed uniformly on $(0, y)$. The marginal distribution of Y (that is, the distribution you would give it if you didn't know the time of the phone company call) is uniform on $(0, 10)$, meaning you have no idea when the shutdown occurred between when you left and when you returned.

(a) What is the posterior pdf on Y given X, the time of the phone company call?

(b) What is a minimum-length 95% posterior probability interval for Y for the observed $X = 1$?

(c) Here are two random intervals:

$$\left\{ y \,\middle|\, y > \frac{X}{.95} \text{ and } y < 10 \right\} \qquad \left\{ y \,\middle|\, y < \frac{X}{.05} \text{ and } y \in (X, 10) \right\} .$$

Show that these are both 95% non-Bayesian confidence sets for Y.

(d) Compare the behavior of these non-Bayesian intervals to that of the Bayesian interval, particularly for X near 0 or near 10.

(e) Food for thought: Can you think of a way to produce better-behaved non-Bayesian confidence intervals here? (Bad behavior, for example: producing empty confidence intervals or 95% intervals that contain the true value of Y with post-sample probability 1.)

(3) Suppose our model is

$$\underset{10\times 1}{Y} = \underset{10\times 1}{X}\beta + \varepsilon$$

$$\varepsilon \mid \{X, \beta, \sigma\} \quad \sim \quad N(0, \sigma^2 I)\,.$$

We have a conjugate prior pdf given by

$$p(\beta, \sigma) = \begin{cases} .5\varphi(\beta; 100\sigma^2) g(\sigma^2)\, d\beta\, d\sigma^2 & \beta \neq 0 \\ .5 g(\sigma^2) & \beta = 0\,, \end{cases}$$

where $g(\sigma^2) = \sigma^{-2} e^{-1/\sigma^2}$ is the inverse-gamma density with parameters 1,1 and $\varphi(\cdot, a)$ is the standard normal pdf with variance a. The prior density is a density with respect to the measure that is Lebesgue measure over $\sigma > 0$ on the subspace of $\mathbb{R}^2$ where $\beta = 0$ and Lebesgue measure over β, σ elsewhere in the parameter space (which is of course the part of $\mathbb{R}^2$ on which $\sigma > 0$). Supposing our sample produces

$$\hat{\beta} = 1,\ \hat{\sigma}^2 = 1,\ X'X = 2\,.$$

Here $\hat{\beta}$ is the OLS estimator and $\hat{\sigma}^2$ is the sum of squared OLS residuals divided by degrees of freedom, which is here 9. What is the posterior probability of $\beta = 0$? What is the smallest significance level at which $H_0 : \beta = 0$ would be rejected by the usual t test that rejects when

$$\frac{\hat{\beta}}{\sqrt{\hat{\sigma}^2 / X'X}}$$

exceeds a critical value?