In this MidTerm there are three questions:

- Q1: 20 Marks
- Q2: 20 Marks
- Q3: 10 Marks

Q1. Consider realization $Y_1, \ldots, Y_n$ from an infinitely exchangeable binary sequence where, in de Finetti's representation, prior distribution π_0 on the unknown parameter θ is specified. Suppose that

$$S_n = \sum_{i=1}^{n} Y_i.$$

Let $y_1, \ldots, y_n$ and s_n be the corresponding observed values. As usual, the posterior distribution will be denoted π_n.

(a) Find $\pi_n(\theta)$ if π_0 is the $Uniform(0, 1)$ distribution. 6 Marks

(b) Find $\pi_n(\theta)$ if π_0 places equal probability on the values $1/3$ and $2/3$ and places zero probability elsewhere. 6 Marks

(c) Suppose that only the value of s_n is recorded. Is it possible to compute π_n given only this information ? Justify your answer. 4 Marks

(d) Find the Bayesian estimate of θ, $\widehat{\theta}_B$, if π_0 is the $Beta(10, 3)$ distribution, and the loss function

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

is used. Recall that

$$\widehat{\theta}_B = \arg\min_t \int L(t, \theta)\pi_n(\theta)\, d\theta.$$

4 Marks

Solution:

(a) We have that the likelihood is

$$\mathcal{L}_n(\theta) = \prod_{i=1}^{n} \theta^{y_i}(1-\theta)^{1-y_i} = \theta^{s_n}(1-\theta)^{n-s_n}$$

and with a uniform prior, this means that

$$\pi_n(\theta) \propto \theta^{s_n}(1-\theta)^{n-s_n} \equiv Beta(s_n+1, n-s_n+1)$$

6 Marks

(b) In this case, the posterior must also concentrate on $\{1/3, 2/3\}$ and therefore the posterior is

$$\pi_n(\theta) \propto \frac{1}{2}\left(\frac{1}{3}\right)^{s_n}\left(\frac{2}{3}\right)^{n-s_n}\mathbb{1}_{\{1/3\}}(\theta) + \frac{1}{2}\left(\frac{2}{3}\right)^{s_n}\left(\frac{1}{3}\right)^{n-s_n}\mathbb{1}_{\{2/3\}}(\theta)$$

that is

$$\pi_n(\theta) = \frac{2^{n-s_n}\mathbb{1}_{\{1/3\}}(\theta) + 2^{s_n}\mathbb{1}_{\{2/3\}}(\theta)}{2^{n-s_n}+2^{s_n}} = \begin{cases} \dfrac{2^{n-s_n}}{2^{n-s_n}+2^{s_n}} & \theta = 1/3 \\ \dfrac{2^{s_n}}{2^{n-s_n}+2^{s_n}} & \theta = 2/3 \end{cases}$$

and zero otherwise. 6 Marks

(c) Yes; from the likelihood above, it is clear that the posterior will depend on the data only through s_n, that is, S_n is a sufficient statistic. 4 Marks

(d) If π_0 is the $Beta(10, 3)$ distribution, we have that

$$\begin{aligned}
\widehat{\theta}_B &= \arg\min_t \int_0^1 \frac{(t-\theta)^2}{\theta(1-\theta)}\pi_n(\theta)\, d\theta \\
&= \arg\min_t \int_0^1 \frac{(t-\theta)^2}{\theta(1-\theta)}\, \theta^{(s_n+10)-1}(1-\theta)^{(n-s_n+3)-1}\, d\theta \quad \text{ignoring multiplicative constants} \\
&= \arg\min_t \int_0^1 (t-\theta)^2\, \theta^{(s_n+9)-1}(1-\theta)^{(n-s_n+2)-1}\, d\theta
\end{aligned}$$

The optimal t therefore corresponds to the mean of the $Beta(s_n+9, n-s_n+2)$ distribution, that is

$$\frac{s_n+9}{n+11}$$

4 Marks

Q2. A Bayesian analysis based on the conditional model

$$f_Y(y;\theta) = \frac{e^{-\theta}\theta^y}{y!} \qquad y = 0, 1, 2, \ldots$$

and zero otherwise, for $\theta > 0$, is to be carried out based on a random sample of size n.

(a) Find a conjugate prior for this problem. 4 Marks

(b) Find the Jeffreys prior for this problem. 6 Marks

(c) Find a large sample approximation to the posterior distribution π_n. 6 Marks

(d) If

$$\pi_0(\theta) = e^{-\theta} \qquad \theta > 0$$

find the posterior distribution for the data set

$$3, 0, 2, 2, 4, 2$$

which represent a sample of size $n = 6$. 4 Marks

Solution:

(a) The likelihood is

$$\mathcal{L}_n(\theta) = \prod_{i=1}^{n} \frac{e^{-\theta}\theta^{y_i}}{y_i!} \propto \theta^{s_n} e^{-n\theta}$$

where $s_n = y_1 + \cdots + y_n$. Therefore a conjugate prior would be

$$\pi_0(\theta) \propto \theta^a e^{-b\theta} \qquad a, b > 0$$

that is $\pi_0(\theta) \equiv Gamma(\alpha_0, \beta_0)$ say.

4 Marks

(b) Here we have, in the usual notation,

$$\ell(y;\theta) = -\theta + y\log\theta + \log y_i!$$
$$\dot{\ell}(y;\theta) = -1 + y/\theta$$
$$\ddot{\ell}(y;\theta) = -y/\theta^2$$

So therefore the Jeffreys prior is

$$\pi_0(\theta) \propto \{\mathbb{E}_Y[Y/\theta^2;\theta]\}^{1/2} = \{1/\theta\}^{1/2}$$

as $\mathbb{E}_Y[Y;\theta] = \theta$. 6 Marks

(c) Using the likelihood theory, we have that

$$\widehat{\theta}_n = \frac{s_n}{n} \qquad \ddot{\ell}_n(\theta) = -s_n/\theta^2$$

so therefore

$$\pi_n(\theta) \doteq Normal(\widehat{\theta}_n, \{-\ddot{\ell}_n(\widehat{\theta}_n)\}^{-1}) \equiv Normal(\widehat{\theta}_n, \widehat{\theta}_n^2/s_n) \equiv Normal(\widehat{\theta}_n, \widehat{\theta}_n/n)$$

6 Marks

(d) Here $s_n = 13$, and therefore

$$\pi_n(\theta) \propto \theta^{13} e^{-6\theta} \times e^{-\theta} = \theta^{14-1} e^{-7\theta} \equiv Gamma(14, 7)$$

4 Marks

Q3. A Bayesian analysis based on the conditional model

$$f_Y(y;\theta) \equiv Normal(\theta, 4)$$

is to be carried out based on a random sample of size n. The conjugate prior

$$\pi_0(\theta) \equiv Normal(\eta_0, 1/\lambda_0)$$

is to be used, for hyperparameters η_0, λ_0.

(a) Find the posterior distribution for θ using the complete-the-square identity

$$A(x-a)^2 + B(x-b)^2 = (A+B)(x-m)^2 + \frac{AB}{A+B}(a-b)^2$$

where

$$m = \frac{Aa+Bb}{A+B}.$$

6 Marks

(b) The Bayes estimate for θ under quadratic loss is the posterior mean. Under what (limiting) condition does the Bayes estimate equal the maximum likelihood estimate ? Justify your answer. 4 Marks

Solution:

(a) Here we need to simplify the quadratic form

$$\frac{1}{4}\sum_{i=1}^{n}(y_i-\theta)^2+\lambda_0(\theta-\eta_0)^2=\frac{1}{4}\sum_{i=1}^{n}(y_i-\overline{y})^2+\frac{n}{4}(\overline{y}-\theta)^2+\lambda_0(\theta-\eta_0)^2$$

which, as a function of θ, we may write as

$$M_n^{-1}(\theta-m_n)^2+\text{constant}$$

where

$$m_n=\frac{n\overline{y}/4+\eta_0\lambda_0}{n/4+\lambda_0}\qquad M_n=\frac{1}{n/4+\lambda_0}$$

Then $\pi_n(\theta)\equiv Normal(m,M)$.

6 Marks

(b) We can see that $\lambda_0\longrightarrow 0$ renders

$$m_n=\overline{y}$$

to be the mle, that is, the value that minimizes the likelihood component of the quadratic form

$$\frac{1}{4}\sum_{i=1}^{n}(y_i-\overline{y})^2+\frac{n}{4}(\overline{y}-\theta)^2$$

with respect to θ. 4 Marks