A1/12 B1/15 **Principles of Statistics**

(i) What does it mean to say that a family $\{f(\cdot|\theta) : \theta \in \Theta\}$ of densities is an *exponential family?*

Consider the family of densities on $(0, \infty)$ parametrised by the positive parameters a, b and defined by

$$f(x|a,b) = \frac{a \exp(-(a-bx)^2/2x)}{\sqrt{2\pi x^3}} \qquad (x > 0).$$

Prove that this family is an exponential family, and identify the natural parameters and the reference measure.

(ii) Let $(X_1, \ldots, X_n)$ be a sample drawn from the above distribution. Find the maximum-likelihood estimators of the parameters (a, b). Find the Fisher information matrix of the family (in terms of the natural parameters). Briefly explain the significance of the Fisher information matrix in relation to unbiased estimation. Compute the mean of X_1 and of X_1^{-1}.

A2/11 B2/16 **Principles of Statistics**

(i) In the context of a decision-theoretic approach to statistics, what is a *loss function*? a *decision rule*? the *risk function* of a decision rule? the *Bayes risk* of a decision rule? the *Bayes rule* with respect to a given prior distribution?

Show how the Bayes rule with respect to a given prior distribution is computed.

(ii) A sample of n people is to be tested for the presence of a certain condition. A single real-valued observation is made on each one; this observation comes from density f_0 if the condition is absent, and from density f_1 if the condition is present. Suppose $\theta_i = 0$ if the i^{th} person does not have the condition, $\theta_i = 1$ otherwise, and suppose that the prior distribution for the θ_i is that they are independent with common distribution $P(\theta_i = 1) = p \in (0,1)$, where p is known. If X_i denotes the observation made on the i^{th} person, what is the posterior distribution of the θ_i?

Now suppose that the loss function is defined by

$$L_0(\theta, a) \equiv \sum_{j=1}^{n} (\alpha a_j (1 - \theta_j) + \beta (1 - a_j) \theta_j)$$

for action $a \in [0,1]^n$, where α, β are positive constants. If π_j denotes the posterior probability that $\theta_j = 1$ given the data, prove that the Bayes rule for this prior and this loss function is to take $a_j = 1$ if π_j exceeds the threshold value $\alpha/(\alpha + \beta)$, and otherwise to take $a_j = 0$.

In an attempt to control the proportion of false positives, it is proposed to use a different loss function, namely,

$$L_1(\theta, a) \equiv L_0(\theta, a) + \gamma I_{\{\sum a_j > 0\}} \left(1 - \frac{\sum \theta_j a_j}{\sum a_j} \right),$$

where $\gamma > 0$. Prove that the Bayes rule is once again a threshold rule, that is, we take action $a_j = 1$ if and only if $\pi_j > \lambda$, and determine λ as fully as you can.

A3/12 B3/15 **Principles of Statistics**

(i) What is a *sufficient statistic*? What is a *minimal* sufficient statistic? Explain the terms *nuisance parameter* and *ancillary statistic*.

(ii) Let $U_1, \ldots, U_n$ be independent random variables with common uniform($[0,1]$) distribution, and suppose you observe $X_i \equiv aU_i^{-\beta}$, $i = 1, \ldots, n$, where the positive parameters a, β are unknown. Write down the joint density of $X_1, \ldots, X_n$ and prove that the statistic

$$(m, p) \equiv (\min_{1 \leqslant j \leqslant n} \{X_j\}, \prod_{j=1}^{n} X_j)$$

is minimal sufficient for (a, β). Find the maximum-likelihood estimator $(\hat{a}, \hat{\beta})$ of (a, β).

Regarding β as the parameter of interest and a as the nuisance parameter, is m ancillary? Find the mean and variance of $\hat{\beta}$. Hence find an unbiased estimator of β.

A4/13 B4/15 **Principles of Statistics**

Suppose that $\theta \in \mathbb{R}^d$ is the parameter of a non-degenerate exponential family. Derive the asymptotic distribution of the maximum-likelihood estimator $\hat{\theta}_n$ of θ based on a sample of size n. [*You may assume that the density is infinitely differentiable with respect to the parameter, and that differentiation with respect to the parameter commutes with integration.*]