Paper 3, Section II

27I Principles of Statistics

What is meant by an *equaliser* decision rule? What is meant by an *extended Bayes* rule? Show that a decision rule that is both an equaliser rule and extended Bayes is minimax.

Let $X_1, \ldots, X_n$ be independent and identically distributed random variables with the normal distribution $\mathcal{N}(\theta, h^{-1})$, and let $k > 0$. It is desired to estimate θ with loss function $L(\theta, a) = 1 - \exp\{-\frac{1}{2}k(a-\theta)^2\}$.

Suppose the prior distribution is $\theta \sim \mathcal{N}(m_0, h_0^{-1})$. Find the *Bayes act* and the *Bayes loss* posterior to observing $X_1 = x_1, \ldots, X_n = x_n$. What is the *Bayes risk* of the Bayes rule with respect to this prior distribution?

Show that the rule that estimates θ by $\overline{X} = n^{-1}\sum_{i=1}^n X_i$ is minimax.

Paper 4, Section II

27I Principles of Statistics

Consider the double dichotomy, where the loss is 0 for a correct decision and 1 for an incorrect decision. Describe the form of a *Bayes decision rule*. Assuming the equivalence of normal and extensive form analyses, deduce the *Neyman–Pearson lemma*.

For a problem with random variable X and real parameter θ, define *monotone likelihood ratio* (MLR) and *monotone test*.

Suppose the problem has MLR in a real statistic $T = t(X)$. Let ϕ be a monotone test, with power function $\gamma(\cdot)$, and let ϕ' be any other test, with power function $\gamma'(\cdot)$. Show that if $\theta_1 > \theta_0$ and $\gamma(\theta_0) > \gamma'(\theta_0)$, then $\gamma(\theta_1) > \gamma'(\theta_1)$. Deduce that there exists $\theta^* \in [-\infty, \infty]$ such that $\gamma(\theta) \leqslant \gamma'(\theta)$ for $\theta < \theta^*$, and $\gamma(\theta) \geqslant \gamma'(\theta)$ for $\theta > \theta^*$.

For an arbitrary prior distribution Π with density $\pi(\cdot)$, and an arbitrary value θ^*, show that the posterior odds

$$\frac{\Pi(\theta > \theta^* \mid X = x)}{\Pi(\theta \leqslant \theta^* \mid X = x)}$$

is a non-decreasing function of $t(x)$.

Paper 1, Section II

28I Principles of Statistics

(i) Let $X_1, \ldots, X_n$ be independent and identically distributed random variables, having the exponential distribution $\mathcal{E}(\lambda)$ with density $p(x|\lambda) = \lambda \exp(-\lambda x)$ for $x, \lambda > 0$. Show that $T_n = \sum_{i=1}^n X_i$ is *minimal sufficient* and *complete* for λ.

[You may assume uniqueness of Laplace transforms.]

(ii) For given $x > 0$, it is desired to estimate the quantity $\phi = \text{Prob}(X_1 > x|\lambda)$. Compute the Fisher information for ϕ.

(iii) State the Lehmann–Scheffé theorem. Show that the estimator $\tilde{\phi}_n$ of ϕ defined by

$$\tilde{\phi}_n = \begin{cases} 0, & \text{if } T_n < x, \\ \left(1 - \frac{x}{T_n}\right)^{n-1}, & \text{if } T_n \geqslant x \end{cases}$$

is the minimum variance unbiased estimator of ϕ based on $(X_1, \ldots, X_n)$. Without doing any computations, state whether or not the variance of $\tilde{\phi}_n$ achieves the Cramér–Rao lower bound, justifying your answer briefly.

Let $k \leqslant n$. Show that $\mathbb{E}(\tilde{\phi}_k \mid T_n, \lambda) = \tilde{\phi}_n$.

Paper 2, Section II

28I Principles of Statistics

Suppose that the random vector $\mathbf{X} = (X_1, \ldots, X_n)$ has a distribution over $\mathbb{R}^n$ depending on a real parameter θ, with everywhere positive density function $p(\mathbf{x} \mid \theta)$. Define the *maximum likelihood estimator* $\hat{\theta}$, the *score variable* U, the *observed information* $\hat{\jmath}$ and the *expected (Fisher) information* I for the problem of estimating θ from $\mathbf{X}$.

For the case where the (X_i) are independent and identically distributed, show that, as $n \to \infty$, $I^{-1/2}\, U \xrightarrow{d} \mathcal{N}(0,1)$. [You may assume sufficient conditions to allow interchange of integration over the sample space and differentiation with respect to the parameter.] State the asymptotic distribution of $\hat{\theta}$.

The random vector $\mathbf{X} = (X_1, \ldots, X_n)$ is generated according to the rule

$$X_{i+1} = \theta X_i + E_i,$$

where $X_0 = 1$ and the (E_i) are independent and identically distributed from the standard normal distribution $\mathcal{N}(0,1)$. Write down the likelihood function for θ based on data $\mathbf{x} = (x_1, \ldots, x_n)$, find $\hat{\theta}$ and $\hat{\jmath}$ and show that the pair $(\hat{\theta}, \hat{\jmath})$ forms a *minimal sufficient statistic.*

A Bayesian uses the improper prior density $\pi(\theta) \propto 1$. Show that, in the posterior, $S(\theta - \hat{\theta})$ (where S is a statistic that you should identify) has the same distribution as E_1.