1/II/27I **Principles of Statistics**

Suppose that X has density $f(\cdot|\theta)$ where $\theta \in \Theta$. What does it mean to say that statistic $T \equiv T(X)$ is *sufficient* for θ?

Suppose that $\theta = (\psi, \lambda)$, where ψ is the parameter of interest, and λ is a nuisance parameter, and that the sufficient statistic T has the form $T = (C, S)$. What does it mean to say that the statistic S is *ancillary*? If it is, how (according to the conditionality principle) do we test hypotheses on ψ? Assuming that the set of possible values for X is discrete, show that S is ancillary if and only if the density (probability mass function) $f(x|\psi, \lambda)$ factorises as

$$f(x|\psi,\lambda) = \varphi_0(x)\ \varphi_C(C(x), S(x), \psi)\ \varphi_S(S(x), \lambda) \qquad (*)$$

for some functions φ_0, φ_C, and φ_S with the properties

$$\sum_{x \in C^{-1}(c) \cap S^{-1}(s)} \varphi_0(x) = 1 = \sum_s \varphi_S(s,\lambda) = \sum_s \sum_c \varphi_C(c,s,\psi)$$

for all c, s, ψ, and λ.

Suppose now that $X_1, \ldots, X_n$ are independent observations from a $\Gamma(a, b)$ distribution, with density

$$f(x|a,b) = (bx)^{a-1} e^{-bx} b I_{\{x>0\}} / \Gamma(a).$$

Assuming that the criterion $(*)$ holds also for observations which are not discrete, show that it is not possible to find $(C(X), S(X))$ sufficient for (a, b) such that S is ancillary when b is regarded as a nuisance parameter, and a is the parameter of interest.

2/II/27I **Principles of Statistics**

(i) State Wilks' likelihood ratio test of the null hypothesis $H_0 : \theta \in \Theta_0$ against the alternative $H_1 : \theta \in \Theta_1$, where $\Theta_0 \subset \Theta_1$. Explain when this test may be used.

(ii) Independent identically-distributed observations $X_1, \ldots, X_n$ take values in the set $S = \{1, \ldots, K\}$, with common distribution which under the null hypothesis is of the form

$$P(X_1 = k|\theta) = f(k|\theta) \qquad (k \in S)$$

for some $\theta \in \Theta_0$, where Θ_0 is an open subset of some Euclidean space $\mathbb{R}^d$, $d < K - 1$. Under the alternative hypothesis, the probability mass function of the X_i is unrestricted.

Assuming sufficient regularity conditions on f to guarantee the existence and uniqueness of a maximum-likelihood estimator $\hat{\theta}_n(X_1, \ldots, X_n)$ of θ for each n, show that for large n the Wilks' likelihood ratio test statistic is approximately of the form

$$\sum_{j=1}^{K} (N_j - n\hat{\pi}_j)^2 / N_j,$$

where $N_j = \sum_{i=1}^{n} I_{\{X_i = j\}}$, and $\hat{\pi}_j = f(j|\hat{\theta}_n)$. What is the asymptotic distribution of this statistic?

3/II/26I **Principles of Statistics**

(i) In the context of decision theory, what is a *Bayes rule* with respect to a given loss function and prior? What is an *extended Bayes rule*?

Characterise the Bayes rule with respect to a given prior in terms of the posterior distribution for the parameter given the observation. When $\Theta = \mathcal{A} = \mathbb{R}^d$ for some d, and the loss function is $L(\theta, a) = \|\theta - a\|^2$, what is the Bayes rule?

(ii) Suppose that $\mathcal{A} = \Theta = \mathbb{R}$, with loss function $L(\theta, a) = (\theta - a)^2$ and suppose further that under P_θ, $X \sim N(\theta, 1)$.

Supposing that a $N(0, \tau^{-1})$ prior is taken over θ, compute the Bayes risk of the decision rule $d_\lambda(X) = \lambda X$. Find the posterior distribution of θ given X, and confirm that its mean is of the form $d_\lambda(X)$ for some value of λ which you should identify. Hence show that the decision rule d_1 is an extended Bayes rule.

4/II/27I Principles of Statistics

Assuming sufficient regularity conditions on the likelihood $f(x|\theta)$ for a univariate parameter $\theta \in \Theta$, establish the Cramér–Rao lower bound for the variance of an unbiased estimator of θ.

If $\hat{\theta}(X)$ is an unbiased estimator of θ whose variance attains the Cramér–Rao lower bound for every value of $\theta \in \Theta$, show that the likelihood function is an exponential family.