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

An angler starts fishing at time 0. Fish bite in a Poisson Process of rate Λ per hour, so that, if $\Lambda = \lambda$, the number N_t of fish he catches in the first t hours has the Poisson distribution $\mathcal{P}(\lambda t)$, while T_n, the time in hours until his nth bite, has the Gamma distribution $\Gamma(n, \lambda)$, with density function

$$p(t \mid \lambda) = \frac{\lambda^n}{(n-1)!} t^{n-1} e^{-\lambda t} \quad (t > 0).$$

Bystander B_1 plans to watch for 3 hours, and to record the number N_3 of fish caught. Bystander B_2 plans to observe until the 10th bite, and to record T_{10}, the number of hours until this occurs.

For B_1, show that $\widetilde{\Lambda}_1 := N_3/3$ is an unbiased estimator of Λ whose variance function achieves the Cramér–Rao lower bound.

Find an unbiased estimator of Λ for B_2, of the form $\widetilde{\Lambda}_2 = k/T_{10}$. Does it achieve the Cramér–Rao lower bound? Is it minimum-variance-unbiased? Justify your answers.

In fact, the 10th fish bites after exactly 3 hours. For each of B_1 and B_2, write down the likelihood function for Λ based their observations. What does the *Likelihood Principle* have to say about the inferences to be drawn by B_1 and B_2, and why? Compute the estimates $\widetilde{\lambda}_1$ and $\widetilde{\lambda}_2$ produced by applying $\widetilde{\Lambda}_1$ and $\widetilde{\Lambda}_2$ to the observed data. Does the method of minimum-variance-unbiased estimation respect the Likelihood Principle?

2/II/27I Principles of Statistics

Under hypothesis H_i $(i = 0, 1)$, a real-valued observable X, taking values in $\mathcal{X}$, has density function $p_i(\cdot)$. Define the *Type I error* α and the *Type II error* β of a test $\phi : \mathcal{X} \to [0, 1]$ of the null hypothesis H_0 against the alternative hypothesis H_1. What are the *size* and *power* of the test in terms of α and β?

Show that, for $0 < c < \infty$, ϕ minimises $c\alpha + \beta$ among all possible tests if and only if it satisfies

$$\begin{aligned} p_1(x) > c\, p_0(x) &\Rightarrow \phi(x) = 1, \\ p_1(x) < c\, p_0(x) &\Rightarrow \phi(x) = 0. \end{aligned}$$

What does this imply about the admissibility of such a test?

Given the value θ of a parameter variable $\Theta \in [0, 1)$, the observable X has density function

$$p(x \mid \theta) = \frac{2(x - \theta)}{(1 - \theta)^2} \qquad (\theta \leqslant x \leqslant 1).$$

For fixed $\theta \in (0, 1)$, describe all the likelihood ratio tests of $H_0 : \Theta = 0$ against $H_\theta : \Theta = \theta$.

For fixed $k \in (0, 1)$, let ϕ_k be the test that rejects H_0 if and only if $X \geqslant k$. Is ϕ_k admissible as a test of H_0 against H_θ for every $\theta \in (0, 1)$? Is it uniformly most powerful for its size for testing H_0 against the composite hypothesis $H_1 : \Theta \in (0, 1)$? Is it admissible as a test of H_0 against H_1?

3/II/26I Principles of Statistics

Define the notion of *exponential family (EF)*, and show that, for data arising as a random sample of size n from an exponential family, there exists a sufficient statistic whose dimension stays bounded as $n \to \infty$.

The log-density of a normal distribution $\mathcal{N}(\mu, v)$ can be expressed in the form

$$\log p(x \mid \boldsymbol{\phi}) \;=\; \phi_1\, x + \phi_2\, x^2 - k(\boldsymbol{\phi})$$

where $\boldsymbol{\phi} = (\phi_1, \phi_2)$ is the value of an unknown parameter $\Phi = (\Phi_1, \Phi_2)$. Determine the function k, and the natural parameter-space $\mathbb{F}$. What is the *mean-value parameter* $\mathrm{H} = (\mathrm{H}_1, \mathrm{H}_2)$ in terms of Φ?

Determine the maximum likelihood estimator $\widehat{\Phi}_1$ of Φ_1 based on a random sample $(X_1, \ldots, X_n)$, and give its asymptotic distribution for $n \to \infty$.

How would these answers be affected if the variance of X were known to have value v_0?

4/II/27I Principles of Statistics

Define *sufficient statistic*, and state the factorisation criterion for determining whether a statistic is sufficient. Show that a Bayesian posterior distribution depends on the data only through the value of a sufficient statistic.

Given the value μ of an unknown parameter M, observables $X_1, \ldots, X_n$ are independent and identically distributed with distribution $\mathcal{N}(\mu, 1)$. Show that the statistic $\overline{X} := n^{-1} \sum_{i=1}^{n} X_i$ is sufficient for M.

If the prior distribution is $\mathrm{M} \sim \mathcal{N}(0, \tau^2)$, determine the posterior distribution of M and the predictive distribution of $\overline{X}$.

In fact, there are two hypotheses as to the value of M. Under hypothesis H_0, M takes the known value 0; under H_1, M is unknown, with prior distribution $\mathcal{N}(0, \tau^2)$. Explain why the *Bayes factor* for choosing between H_0 and H_1 depends only on $\overline{X}$, and determine its value for data $X_1 = x_1, \ldots, X_n = x_n$.

The frequentist 5%-level test of H_0 against H_1 rejects H_0 when $|\overline{X}| \geqslant 1.96/\sqrt{n}$. What is the Bayes factor for the critical case $|\overline{x}| = 1.96/\sqrt{n}$? How does this behave as $n \to \infty$? Comment on the similarities or differences in behaviour between the frequentist and Bayesian tests.