Paper 2, Section II

26K Principles of Statistics

We consider the problem of estimating θ in the model $\{f(x,\theta):\ \theta \in (0,\infty)\}$, where

$$f(x,\theta) = (1-\alpha)(x-\theta)^{-\alpha}\mathbf{1}\{x \in [\theta, \theta+1]\}\,.$$

Here $\mathbf{1}\{A\}$ is the indicator of the set A, and $\alpha \in (0,1)$ is known. This estimation is based on a sample of n i.i.d. $X_1,\ldots,X_n$, and we denote by $X_{(1)} < \ldots < X_{(n)}$ the ordered sample.

(a) Compute the mean and the variance of X_1. Construct an unbiased estimator of θ taking the form $\tilde{\theta}_n = \bar{X}_n + c(\alpha)$, where $\bar{X}_n = n^{-1}\sum_{i=1}^n X_i$, specifying $c(\alpha)$.

(b) Show that $\tilde{\theta}_n$ is consistent and find the limit in distribution of $\sqrt{n}(\tilde{\theta}_n - \theta)$. Justify your answer, citing theorems that you use.

(c) Find the maximum likelihood estimator $\hat{\theta}_n$ of θ. Compute $\mathbf{P}(\hat{\theta}_n - \theta > t)$ for all real t. Is $\hat{\theta}_n$ unbiased?

(d) For $t > 0$, show that $\mathbf{P}(n^{\beta}(\hat{\theta}_n - \theta) > t)$ has a limit in $(0,1)$ for some $\beta > 0$. Give explicitly the value of β and the limit. Why should one favour using $\hat{\theta}_n$ over $\tilde{\theta}_n$?

Paper 3, Section II

26K Principles of Statistics

We consider the problem of estimating an unknown θ_0 in a statistical model $\{f(x,\theta)\,,\,\theta \in \Theta\}$ where $\Theta \subset \mathbb{R}$, based on n i.i.d. observations $X_1, \ldots, X_n$ whose distribution has p.d.f. $f(x,\theta_0)$.

In all the parts below you may assume that the model satisfies necessary regularity conditions.

(a) Define the *score function* S_n of θ. Prove that $S_n(\theta_0)$ has mean 0.

(b) Define the *Fisher Information* $I(\theta)$. Show that it can also be expressed as

$$I(\theta) = -\mathbb{E}_\theta\Big[\frac{d^2}{d\theta^2}\log f(X_1,\theta)\Big]\,.$$

(c) Define the *maximum likelihood estimator* $\hat{\theta}_n$ of θ. Give without proof the limits of $\hat{\theta}_n$ and of $\sqrt{n}(\hat{\theta}_n - \theta_0)$ (in a manner which you should specify). [Be as precise as possible when describing a distribution.]

(d) Let $\psi : \Theta \to \mathbb{R}$ be a continuously differentiable function, and $\tilde{\theta}_n$ another estimator of θ_0 such that $|\hat{\theta}_n - \tilde{\theta}_n| \leqslant 1/n$ with probability 1. Give the limits of $\psi(\tilde{\theta}_n)$ and of $\sqrt{n}(\psi(\tilde{\theta}_n) - \psi(\theta_0))$ (in a manner which you should specify).

Paper 4, Section II

27K Principles of Statistics

For the statistical model $\{\mathcal{N}_d(\theta, \Sigma), \theta \in \mathbb{R}^d\}$, where Σ is a known, positive-definite $d \times d$ matrix, we want to estimate θ based on n i.i.d. observations $X_1, \ldots, X_n$ with distribution $\mathcal{N}_d(\theta, \Sigma)$.

(a) Derive the maximum likelihood estimator $\hat{\theta}_n$ of θ. What is the distribution of $\hat{\theta}_n$?

(b) For $\alpha \in (0,1)$, construct a confidence region C_n^α such that $\mathbf{P}_\theta(\theta \in C_n^\alpha) = 1 - \alpha$.

(c) For $\Sigma = I_d$, compute the maximum likelihood estimator of θ for the following parameter spaces:

(i) $\Theta = \{\theta : \|\theta\|_2 = 1\}$.

(ii) $\Theta = \{\theta : v^\top \theta = 0\}$ for some unit vector $v \in \mathbb{R}^d$.

(d) For $\Sigma = I_d$, we want to test the null hypothesis $\Theta_0 = \{0\}$ (i.e. $\theta = 0$) against the composite alternative $\Theta_1 = \mathbb{R}^d \setminus \{0\}$. Compute the likelihood ratio statistic $\Lambda(\Theta_1, \Theta_0)$ and give its distribution under the null hypothesis. Compare this result with the statement of Wilks' theorem.

Paper 1, Section II

28K Principles of Statistics

For a positive integer n, we want to estimate the parameter p in the binomial statistical model $\{\text{Bin}(n,p), p \in [0,1]\}$, based on an observation $X \sim \text{Bin}(n,p)$.

(a) Compute the maximum likelihood estimator for p. Show that the posterior distribution for p under a uniform prior on $[0,1]$ is $\text{Beta}(a,b)$, and specify a and b.

[The p.d.f. of $\text{Beta}(a,b)$ is given by

$$f_{a,b}(p) = \frac{(a+b-1)!}{(a-1)!(b-1)!} p^{a-1}(1-p)^{b-1} \,. \,]$$

(b) (i) For a risk function L, define the *risk* of an estimator $\hat{p}$ of p, and the *Bayes risk* under a prior π for p.

(ii) Under the loss function

$$L(\hat{p}, p) = \frac{(\hat{p}-p)^2}{p(1-p)} \,,$$

find a Bayes optimal estimator for the uniform prior. Give its risk as a function of p.

(iii) Give a minimax optimal estimator for the loss function L given above. Justify your answer.