

Exam for **Fundamental Theory of Statistical Inference, 2026.**

You should answer *all* parts, which are equally weighted.

(1) Consider a Bayesian statistical decision problem involving parameter θ , random variable Y and loss function $L(\theta, a)$.

In a situation when no data can be observed, and given a prior probability density function $\pi(\theta)$, define the *Bayes action* to minimise the Bayes loss, $\int L(\theta, a)\pi(\theta)d\theta$. Let $R_0 \equiv R_0(\pi)$ be the Bayes loss of the Bayes action. Express the Bayes risk R_1 of the optimal (Bayes) decision rule when data can be observed in terms of R_0 and the joint distribution of (θ, Y) .

The real-valued parameter θ has prior defined by probability density function $\pi(\theta)$. Consider the problem of specifying a confidence set $S \subseteq \mathbf{R}$ such that the loss under parameter value θ is $L(\theta, S) = c|S| - \mathbf{I}_S(\theta)$, where $\mathbf{I}_S$ is the indicator function of S (so that $\mathbf{I}_S(\theta) = 1$ if $\theta \in S$, $= 0$, otherwise), where $0 < c < 1/\sqrt{2\pi}$ is a specified constant, and where $|S| = \int_S dx$. Show that the ‘highest density’ region $S^* = \{\theta : \pi(\theta) \geq c\}$ provides a Bayes action for this decision problem, and explain why $R_0(\pi) \leq 0$.

For the case where the prior distribution on θ is $N(\mu, \sigma^2)$, find an expression for R_0 in terms of the standard normal distribution function Φ .

Consider the specific case $\mu = 0, \sigma^2 = 1$, and suppose that the distribution of $Y|\theta$ is $N(\theta, 1/9)$. Determine R_1 and show that $R_1 < R_0$.

(2) Let $Y_1, \dots, Y_n$ be independent, identically distributed normal $N(\mu, \sigma^2)$ random variables.

Discuss this distribution as an example of a full exponential family. Justifying your answer carefully, identify the minimum variance unbiased estimators of μ and σ^2 .

Explain how to test the null hypothesis $H_0 : \sigma \leq \sigma_0$ against the alternative hypothesis $H_1 : \sigma > \sigma_0$, for specified σ_0 , for the cases: (a) μ is *known*; (b) μ is *unknown*. What optimality properties does the test have in each case?

(3) ‘The Conditionality Principle is a necessary component of any reasonable approach to frequentist statistical inference’. Discuss, in no more than 300–400 words.