Exam for **Fundamental Theory of Statistical Inference, 2024.**

You should answer *all* parts, which are of equal weight.

(1) Let $Z_1, \ldots, Z_p$ be independent, with Z_i distributed as $N(\mu_i, 1)$. Then $W = Z_1^2 + \cdots + Z_p^2$ is said to be distributed as 'non-central chi-squared with p degrees of freedom and non-centrality parameter $\lambda = \mu_1^2 + \cdots + \mu_p^2$', denoted $\chi^2(p, \lambda)$.

Letting $\chi^2_{0.5}(p, \lambda)$ denote the median of $\chi^2(p, \lambda)$, Robert (1990) showed that, for all $\lambda \geq 0$,

$$p - 1 + \lambda \leq \chi^2_{0.5}(p, \lambda) \leq \chi^2_{0.5}(p, 0) + \lambda.$$

Let $\widehat{\theta}_1$ and $\widehat{\theta}_2$ be estimators of a $p-$dimensional parameter θ, and suppose the Euclidean loss function $L(\theta, \widehat{\theta}) = (\theta - \widehat{\theta})^T(\theta - \widehat{\theta})$ is specified. Then the probability that $\widehat{\theta}_1$ is closer to θ than $\widehat{\theta}_2$ is

$$PN_\theta(\widehat{\theta}_1, \widehat{\theta}_2) = P_\theta\{L(\theta, \widehat{\theta}_1) < L(\theta, \widehat{\theta}_2)\}.$$

We would say that $\widehat{\theta}_2$ is *inadmissible in terms of closeness* if there exists an estimator $\widehat{\theta}_1$ such that $PN_\theta(\widehat{\theta}_1, \widehat{\theta}_2) > 1/2$, for all θ.

(a) Let $Y = (Y_1, \ldots, Y_p)^T, p \geq 3$, have a $p-$dimensional normal distribution with mean $\mu = (\mu_1, \ldots, \mu_p)^T$ and identity covariance matrix [so that $Y_1, \ldots, Y_p$ are independent and Y_i is distributed as $N(\mu_i, 1)$]. Assuming the Euclidean loss function, as above, consider the class of shrinkage estimators given by

$$\left\{\widehat{\mu}_c : \widehat{\mu}_c = \left(1 - \frac{c}{Y^TY}\right) Y\right\},$$

for constants $c \geq 0$. Then $\widehat{\mu}_0 \equiv Y$ and $\widehat{\mu}_{p-2}$ is the James-Stein estimator.

Let $c_1 > c_2$, with corresponding estimators in the shrinkage class

$$\widehat{\mu}_i = \left(1 - \frac{c_i}{Y^TY}\right) Y, i = 1, 2.$$

Show that $PN_\mu(\widehat{\mu}_1, \widehat{\mu}_2) = P\{\chi^2(p, \lambda) > (c_1 + c_2)/2 + \lambda\}$, for a function λ of μ which you should specify.

(b) Show that the theoretically optimal ['oracle'] value of c would be $c^* = \chi^2_{0.5}(p, \lambda) - \lambda$. Why does this value c^* not provide a useful estimator?

(c) By comparing with the James-Stein estimator, show that Y is inadmissible in terms of closeness. By considering the estimator $\widehat{\mu}_{p-1}$, or otherwise, show that the James-Stein estimator is inadmissible in terms of closeness.

CONTINUED

(d) Consider now the case $\mu = \mathbf{0} = (0, \ldots, 0)^T$, and let $D = L(\mathbf{0}, \widehat{\mu}_{p-2}) - L(\mathbf{0}, Y)$ be the difference in losses of the James-Stein estimator and Y.

Show that $P_0\{D < E_0(D)\} > 0.5$. [The subscript 0 denotes that we are fixing $\mu = 0$.] How does $P_0\{D < E_0(D)\}$ behave as a function of p?

(2) Let $Y_1, \ldots, Y_n$ be independent, identically distributed, with common inverse Gaussian probability density function

$$f(y; \mu, \lambda) = \frac{\sqrt{\lambda}}{\sqrt{2\pi}} y^{-3/2} \exp\left\{ -\frac{\lambda(y-\mu)^2}{2\mu^2 y} \right\}, \; y > 0, \; \lambda > 0, \; \mu > 0.$$

(a) Discuss this distribution as an example of a full exponential family.

(b) Suppose that μ is *known*. What is the minimum variance unbiased estimator of λ? Justify your answer.

(c) Explain in detail how to test the null hypothesis $H_0 : \lambda \leq \lambda_0$ against the alternative hypothesis $H_1 : \lambda > \lambda_0$, in the case that μ is *unknown*. Explain carefully the justification of the test and explain any UMP or UMPU properties it may possess.

[You may note that $\frac{\lambda}{\mu^2}(Y-\mu)^2/Y$ is χ_1^2, chi-squared with one degree of freedom. Further, $V = \sum_{i=1}^n (Y_i^{-1} - \bar{Y}^{-1})$ is independent of $\bar{Y} = n^{-1} \sum_{i=1}^n Y_i$, with λV distributed as χ_{n-1}^2.]

(3) 'The role of unbiasedness in statistical theory is ambiguous'. Discuss, in no more than about 250–300 words.